




DEPARTMENT OF VETERANS AFFAIRS



OFFICE OF INFORMATION AND TECHNOLOGY
SERVICE DELIVERY AND ENGINEERING
ENTERPRISE SYSTEMS ENGINEERING

Red Hat Enterprise Linux 7 Server Baseline

Version 1
September 21, 2015

Revision History

Revision #	Revision Date	Description of Change	Author
0.1	09/17/14	Initial Draft	Donald Adams
0.2	10/20/14	Minor changes	James Owens
0.3	12/1/14	Removed References to DoD	Donald Adams
0.4	12/11/14	Edited for clarity	Donald Adams
0.5	12/12/14	Technical edits	Raven Nuckols
0.6	1/14/15	Edits requested by VA	Donald Adams
0.7	1/15/15	Revisions of tech edits	Raven Nuckols
0.8	6/5/15	Edits requested by VA	Donald Adams
1.0	7/21/2015	Updated by VA	John Dellar

Distribution

[illegible]

Table of Contents

1	Overview	4
1.1	Conventions.....	4
2	Purpose.....	5
3	Policy.....	5
4	Scope of Authority:	6
5	Responsibilities	6
6	Procedures	6
6.1	RHEL7 Baseline Configuration and Security	6
6.1.1	Server Availability Options	6
6.2	Red Hat Enterprise Satellite Server.....	6
6.3	RHEL Required Channel Subscriptions.....	7
6.4	RHEL Required Applications Installation	7
6.5	File system Configuration	8
6.5.1	Partition/Logical Volume Layout	8
6.5.2	Restrict Mount Point options	10
6.5.3	Bind Mount the /var/tmp directory to /tmp.....	11
6.5.4	Set Sticky Bit on all World-Writable Directories	11
6.6	Software Package Updates	12
6.7	Verify that Red Hat GPG Key is installed	12
6.8	Ensure gpgcheck Enabled In Main Yum Configuration	13
6.9	Verify Package Integrity Using RPM	13
6.10	Verify File Permissions Using RPM.....	13
6.11	Advanced Intrusion Detection Environment (AIDE).....	14
6.11.1	Install AIDE	14
6.11.2	Initialize AIDE.....	14
6.11.3	Implement Periodic Execution of File Integrity Check	15
6.12	Disable Unused Services	15
6.12.1	Disable Interactive startup	16

6.12.2	Verify grub.cfg settings	17
6.12.3	Shadow Passwords.....	17
6.12.4	Verify that no Non-Root Accounts have a UID of 0	17
6.12.5	Verify settings on /etc/shadow	17
6.12.6	Verify settings on /etc/gshadow.....	17
6.12.7	Verify settings on /etc/group.....	18
6.12.8	Verify settings on /etc/passwd	18
6.12.9	Ensure Password File Integrity	18
6.12.10	Lock Inactive User Accounts	18
6.13	Account Management	19
6.13.1	Mandatory Password Management Practices	19
6.13.2	Service Account Management	20
6.13.3	Creating Unique Local User Accounts	21
6.13.4	The SUDO command must require authentication.....	23
6.14	RHEL Configuration Channels	23
6.14.1	VA Baseline RHEL 7.....	24
6.14.2	Modify Default System-wide Profiles	25
6.14.3	Disable X Windows startup at boot time	27
6.14.4	Configuring Secure Sockets Layer (SSH) Server	27
6.14.5	Configure /etc/login.defs to meet password default requirements.....	29
6.14.6	Password Protect Single User Mode.....	29
6.14.7	Logging	30
6.14.8	Disable Unused Kernel Drivers	31
6.14.9	System Time Protocol Service (NTP).....	35
6.14.10	Configuring PAM for user account compliance	36
6.14.11	Restrict Root Logins to the System Console.....	38
6.14.12	Restricting use of the SU commands	39
6.14.13	Ensure Password Hashing Algorithm is SHA-512	39
6.14.14	Implement Inactivity Timeout for Login Shells.....	40
6.14.15	Disable ipv6.....	41
6.14.16	System Resource Management	41
6.14.17	Enable Randomized Layout of Virtual Address Space	43
6.14.18	Configure ExecShield	43

6.14.19	Disable Ctrl-Alt-Delete Reboot Activation.....	44
6.14.20	Preventing Network Attacks	44
6.14.21	Configure System Auditing with Auditd	46
6.14.22	Configure SELINUX.....	50
6.14.23	Mail must be forwarded to one or more system administrators	50
6.15	VA V2S RHEL7.....	51
6.15.1	Configure SNMP	51
7	References	53
8	Review Date	53
9	Follow-Up Responsibility	53
Appendix A Functional Configuration Baselines		54
Intel CPU's		54
Issue - C-State limits		54
Dell Systems		55
Issue - Biosdevname.....		55
Virtual Machines		56
Hypervisor specific client tools.....		56

Figures

Tables

Table 1 – Note and Warning Symbols.....	4
Table 2 – Typographic Conventions.....	4

1 OVERVIEW

1.1 CONVENTIONS

[Table 1 – Note and Warning Symbols](#) lists the Note and Warning Symbol conventions used in this guide to alert the reader to special information or conditions. [Table 2 – Typographic Conventions](#) lists the typographic conventions used in the document.

Table 1 – Note and Warning Symbols



Symbol	Note Type	Description
	NOTE:	Emphasizes points, remind readers of something, or to indicate minor problems in the outcome of what they are doing.
	REFERENCE:	Directs the reader to other sources of information.
	DISCLAIMER:	Generally any statement intended to specify or delimit the scope of rights and obligations that may be exercised and enforced by parties in a legally recognized relationship.
	DANGER:	Warns readers about the possibility of serious or fatal injury to themselves or others.
	WARNING:	Alerts readers about the possibility of minor injury to themselves or others.
	CAUTION:	Draws special attention to anything that could damage equipment or cause the loss of data.

Table 2 – Typographic Conventions

Font, Style or Symbol	Represents	Examples:
Blue text, underlined	External hyperlink to another document or Uniform Resource Locator (URL)	<i>For more information, please refer to the Department of Veterans Affairs website.</i>
Green text, dotted underlined	Hyperlink within this document	<i>See Overview for more information.</i>
Arial	Text inside tables	<i>This is an example of the text normally found within a table.</i>
Courier New (Monospace) with gray framed background	Computer code or computer screen text	<pre>inetd.conf > inetd.conf.new cp inetd.conf.new inetd.conf rm -f /etc/inetd.conf.new</pre>
Courier New, bold	Menu options	Task Manager → View Task Schedule → Schedule New Task
	Screen prompts	Copy file? Y/N
	Directories, file names, command names, system responses, command line commands	<i>...using cache.cpf configuration file.</i>

	Field Names	<i>In the Indicator field, enter the logic that is to be used to determine if the test was positive for the selected MDRO.</i>
	Keyboard keys	<F1> , <Alt> , <L> , <Tab> , <Enter>
Courier New, bold, red	Code and deployment descriptor samples to indicate lines of particular interest, discussed in the preceding text	<p>scd equals the site's three letter site code, 999 should be replaced by the site's station number and z corresponds to the number used in the application instance name.</p> <pre>\$@SYS\$MANAGER:CSTOP scd999A0z QUIETLY \$@SYS\$MANAGER:CSTOP scd999SVR QUIETLY</pre>
→	Choosing a command from a cascading menu	Task Manager → View Task Schedule → Schedule New Task
Times New Roman	Body text	<i>There are no changes in the performance of the system once the installation process is complete.</i>
Times New Roman Italic	Text emphasis	<i>It is very important...</i>
	National and International Standard names	<i>International Statistical Classification of Diseases and Related Health Problems</i>
	Document names	<i>Microsoft Manual of Style for Technical Publications</i>

2 PURPOSE

The purpose of this document is to provide guidance for the base installation, configuration and securing of Red Hat Enterprise Linux 7 (RHEL 7), operating system for implementation onto production servers within the Veterans Affairs (VA) environment.

The objectives of this policy:

To produce a secure and less vulnerable network by ensuring the systems connected to the network are correctly updated and secured.

- To standardize the build and images across the Enterprise of all systems running RHEL 7
- This document is to be used for the hardening, and as a Baseline Configuration for securing RHEL version 7 servers.
- To implement National Institute of Standards and Technology (NIST)
- It is not the objective of this policy to address management and reconfiguration actions.

3 POLICY

This policy is the Baseline Configuration Standard for Red Hat Enterprise Linux Servers. It is for the installation and security hardening of Red Hat Enterprise Linux version 7 Operating System (OS) within the VA.

4 SCOPE OF AUTHORITY:

The content and direction of this policy are in accordance with the VA Handbook 6500, Defense Information Systems Administration (DISA) RHEL checklists, and NIST Checklists.

5 RESPONSIBILITIES

Material Weakness Guide Design: Enterprise Platform Engineering (EPE).

6 PROCEDURES

The following section will provide manual detailed procedures to be used in conducting a review of the RHEL 7 systems for Security Readiness Review process. These procedures are based on documents provided by NIST and the Red Hat RHEL 7 Security Guide.

6.1 RHEL7 BASELINE CONFIGURATION AND SECURITY

6.1.1 Server Availability Options

- Thermal Shutdown – Enabled
- Wake On Local Area Network (LAN) – Enabled
- Power-on Self-Test (POST) F1 Prompt – Disabled
- Power Button – Enabled
- Auto Power On – Enabled
- Power on Delay – No Delay
- Boot Order – Boot from Hard Disk first
 - CD ROM second
 - PXE boot third
 - Other options as may be necessary
- Set Basic Input/Output System (BIOS) Password – *Ensure password is secure*
 - (only required if server is not located in a secure location)

6.2 RED HAT ENTERPRISE SATELLITE SERVER

All Production Red Hat Linux systems are required to be registered with and receive updates from the National Satellite Server. In order to ensure consistency with Operating Systems deployment, the National Satellite administrators will establish clones of the base Red Hat channels that will have the prefix *VA-PREPROD* and *VA-PROD* which will be the only channels visible to the regions...

On the first weekend of each month new and updated packages will be added to the VA-PREPROD channels, so that all Regions and other entities can complete their own testing and evaluation.

On the first weekend of each quarter new and updated packages from VA-PREPROD will be placed into VA- PROD after being processed through ESSCB and NCCB approval process.

So that regions are able to comply with the timeframes required for applying patches as outlined in the Memorandum, Patch Management and Compliance VA Intranet Quorum (VAIQ) #7294131) of Nov 16 2012.

- a. Critical-patches will be tested and applied within 30 days
- b. High-patches will be tested and applied within 60 days.
- c. Moderate-patches will be tested and applied within 90 days
- d. Low-system administrators will determine patching timeframe.
- e. Emergent-patches will be applied ASAP

National Satellite Administrators will test critical and if time allows, high level patches, every week. High and moderate patches will be tested within 30 days of release. All others will be moved to the VA-PREPROD channel within 30 days of release.

Note: Not every patch will be tested due to time constraints

6.3 RHEL REQUIRED CHANNEL SUBSCRIPTIONS

A software channel is a collection of Red Hat Package Manager (RPM) packages. These are the packages that are deployed on systems managed by the National Satellite Server. Software channels define which packages a given system has access to. All Red Hat Linux servers are subscribed to at least the following channels. Other channels may be subscribed to as necessary for a particular application or server function.

- prod-rhel-x86_64-server-7
- prod-rhn-tools-rhel-x86_64-server-7
- prod-bigfix_x64
- prod-encase_V2s

Each region will create their own production clones of these national channels

6.4 RHEL REQUIRED APPLICATIONS INSTALLATION

Current approved Visibility to Servers (V2S) suite as approved for Linux Systems. .

- BigFix
- Encase

- Simple Network Management Protocol (SNMP)
- McAfee Anti-Virus Software (See Note Below)
- McAfee EPO agent (See Note Below)
- McAfee Host Intrusion Prevention (HIPS). (See Note Below)

i Note: Before installing McAfee Software products, ensure that the systems that they are being installed on have been tested and are compatible

6.5 FILE SYSTEM CONFIGURATION

Directories that are used for system-wide functions can be protected by placing them on separate partitions. This provides protection from resource exhaustion and enables the use of mounting options that are applicable to the directory's intended use. User's data can be stored on separate partitions and have stricter mount options.

The overall strategy is to apply security concepts to OS partitioning in the beginning of system build and installation activities in such a way that mount options can be applied to secure the machine—such as noexec options for /tmp and /home, etc. Red Hat does not recommend putting partitions, directories and files in non-standard locations, since SELinux already has a well-defined concept of where things belong.

Logical Volume Management (LVM) should be used so as to provide increased robustness and flexibility in disk and partition management.

XFS is the default file system for RHEL 7, however the tools and utilities that are used with EXT4 no longer work with, or work differently on XFS, and there are new tools (XFSPROGS) that are used to work with XFS file systems.

- EXT4 file system may be used to retain compatibility with current tools and configurations.
- XFS can be used if so desired.

6.5.1 Partition/Logical Volume Layout

The following partitioning scheme is the minimum required as it protects and segments various portions of operating system functionality. Such partitioning provides various benefits, among them being the protection of any one of them becoming full not adversely affecting the others.

/ – The root partition is required; it is the top of the UNIX file system tree structure and must be established as its own partition.

/boot – Required partition; allocated on the first physical drive the BIOS will look to for boot and/or start-up information.

/home – Will be its own distinct partition as a repository for local storage of administrative and user files. The underlying issue is that any place a user can write to must be on a separate partition from /bin, /sbin, and /usr.

/tmp – The /tmp partition is a world-writable directory that provides a safe container for temporary storage and lessens the possibility of resource exhaustion should it become full.

/var – The /var partition is a container for libraries and system services to store frequently changing data and for temporary application files.

/var/log – container for application and security logs.

/var/log/audit – Audit data is stored on its own partition so that auditd will correctly calculate when its partition is out of space.

<swap> – provides additional virtual memory to supplement physical memory. Swap size can be determined by using the settings below:

<2GB RAM	2 times the amount of RAM
>2GB – 8GB RAM	Equal to the amount of RAM
>8GB – 64GB RAM	0.5 times the amount of RAM
>64GB RAM	32GB of swap space

Other, additional partitions may be added, and sized, as site/mission/server function requirements dictate. These include /opt, /srv /usr/local and other data partitions.

```
partition /boot --fstype ext4 --size=512
partition pv.2 --size=1 --grow
volgroup vg_root --pesize=32768 pv.2
logvol swap --fstype swap --name=lv_swap --vgname=vg_root --size=16000
logvol / --fstype ext4 --name=lv_root --vgname=vg_root --fsoptions="noatime" --
size=12000
logvol /home --fstype ext4 --name=lv_home --vgname=vg_root --fsoptions="noatime,nodev"
--size=4096
logvol /var --fstype ext4 --name=lv_var --vgname=vg_root --fsoptions="noatime,nodev" -
size=8000
logvol /var/log --fstype ext4 --name=lv_logs --vgname=vg_root --
fsoptions="noatime,nodev" --size=4096
logvol /var/log/audit --fstype ext4 --name=lv_audit --vgname=vg_root --
fsoptions="noatime,nodev" --size=2000
logvol /tmp --fstype ext4 --name=lv_tmp --vgname=vg_root --
fsoptions="noatime,nodev,nosuid" --size=4096
logvol /srv --fstype ext4 --name=lv_srv --vgname=vg_root --size=1 --grow --percent=25
--fsoptions="noatime,nodev" 2
```

Example partitioning scheme

/dev/mapper/vg_root-lv_root	/	ext4 defaults,noatime	1 2
-----------------------------	---	-----------------------	-----

```

/dev/mapper/vg_root-lv_var    /var      ext4 defaults,noatime,nodev 1 2
/dev/mapper/vg_root-lv_tmp    /tmp      ext4 defaults,noatime,nodev,nosuid,noexec 1 2
/dev/mapper/vg_root-lv_logs   /var/log  ext4 defaults,noatime,nodev 1 2
/dev/mapper/vg_root-lv_audit  /var/log/audit ext4 defaults,noatime,nodev 1 2
/dev/mapper/vg_root-lv_home   /home     ext4 defaults,noatime,nodev 1 2
/dev/mapper/dataavg1-u01      /u01      ext4 defaults,noatime,nodev 1 2
/dev/mapper/vg_root-swap      swap       swap  defaults 0 0

```

Example /etc/fstab

i Note: The *noatime* setting is not a security setting, but is a performance booster that prevents the file system recording each time a file is read, but does not stop the time that a file is modified from being recorded.

6.5.2 Restrict Mount Point options

System partitions can be mounted with various options which limit what files on those partitions can do. These options are set in the file */etc/fstab*, and can be used to make certain types of malicious behavior more difficult.

These settings can be set at installation time as part of the partitioning scheme by using the *–fsoptions* parameter

6.5.2.1 Set nodev option to Non-Root Local Partitions

The *nodev* option prevents users from mounting unauthorized devices on any partition or logical volume which is known not to contain any authorized devices. These include:

- /var
- /var/log/
- /var/log/audit
- /home
- /tmp/
- /dev/shm

The root partition typically contains the */dev* partition, which is the primary location for authorized devices, so this option should not be set on the root / partition .

6.5.2.2 nosuid option

The nosuid mount option specifies that the file system does not allow the setuid or setgid bits to take effect.

- /tmp
- /dev/shm

6.5.2.3 Set *noexec* option

The noexec mount option specifies that although the file system may contain executable scripts and binaries, the OS is not allowed direct execution of those files.

- /tmp
- /dev/shm

During the installation of some applications, the install will fail as the installer tries to extract and execute temporary scripts to */tmp*. To temporarily disable the noexec option run the follow commands

```
mount -o remount,exec /tmp
#Once the application has been installed, re-enalbe the default mount options
specified in fstab with:
mount -o remount /tmp
```

6.5.3 Bind Mount the /var/tmp directory to /tmp

The /var/tmp directory is normally a standard subdirectory in the /var directory. Binding /var/tmp to /tmp establishes an unbreakable link that cannot be removed (even by the root user). It also allows /var/tmp to inherit the same mount options that /tmp owns, allowing /var/tmp to be protected in the same way that /tmp is protected. It will also prevent /var from filling up with temporary files as the contents of /var/tmp will actually reside in the file system containing /tmp. This will also prevent a user from running the /var file system out of space or trying to perform operations that have been blocked in the /tmp file system.

Edit the */etc/fstab* file to and add or modify the following line:

```
/tmp /var/tmp none bind 0 0
```

6.5.4 Set Sticky Bit on all World-Writable Directories

Setting the sticky bit on world writable directories prevents users from deleting or renaming files in that directory that are not owned by them.

To find and remediate incorrectly set directories, run the following command:

```
df --local -P | awk {'if (NR!=1) print $6'} \
| xargs -I '{}' find '{}' -xdev -type d \
\(-perm -0002 -a ! -perm -1000\) 2>/dev/null\
| xargs chmod a+t
```

To find directories that do not have the sticky bit set, run the following command:

```
df --local -P | awk {'if (NR!=1) print $6'} \
| xargs -I '{}' find '{}' -xdev -type d \
\(-perm -0002 -a ! -perm -1000\) 2>/dev/null
```

6.6 SOFTWARE PACKAGE UPDATES

Operating systems and application software need to routinely have patches and updates applied on a regular basis. Updates should be applied as often as possible, with a maximum time between updates of 90 days.

All Production Red Hat Linux systems are to be registered and updates applied from the national Red Hat Satellite. Each region has a Satellite Proxy Server that is used to reduce the network load from the National Satellite. The system administrator will be responsible for installing all patches and package updates.

6.7 VERIFY THAT RED HAT GPG KEY IS INSTALLED

Red Hat cryptographically signs updates with a GNU Privacy Guard (GPG) key to verify that they are valid.

It is important to ensure that updates are obtained from a valid source to protect against spoofing that could lead to the inadvertent installation of malware on the system.

The following command can be used to print the installed release key's fingerprint, which is actually contained in the file referenced below:

```
# gpg --quiet --with-fingerprint /etc/pki/rpm-gpg/RPM-GPG-KEY-redhat-release

pub 4096R/FD431D51 2009-10-22 Red Hat, Inc. (release key 2) <security@redhat.com>
   Key fingerprint = 567E 347A D004 4ADE 55BA 8A5F 199E 2F91 FD43 1D51
pub 1024D/2FA658E0 2006-12-01 Red Hat, Inc. (auxiliary key) <security@redhat.com>
```

```
Key fingerprint = 43A6 E49C 4A38 F4BE 9ABF 2A53 4568 9C88 2FA6 58E
```

Run the following command to ensure that the system has the Red Hat GPG key properly installed:

```
# rpm -q --queryformat "%{SUMMARY}\n" gpg-pubkey  
  
gpg(Red Hat, Inc. (release key 2) <security@redhat.com>)  
gpg(Red Hat, Inc. (auxiliary key) <security@redhat.com>)
```

6.8 ENSURE GPGCHECK ENABLED IN MAIN YUM CONFIGURATION

The *gpgcheck* option should be used to ensure checking of an RPM package's signature always occurs prior to its installation. To configure yum to check package signatures before installing them, ensure the following line appears in */etc/yum.conf* in the [main] section. This is the default in Red Hat 7.

```
grep gpgcheck /etc/yum.conf  
gpgcheck=1
```

6.9 VERIFY PACKAGE INTEGRITY USING RPM

RPM has the capability of verifying installed packages by comparing the installed files against the file information stored in the package. Verifying packages gives a system administrator the ability to detect if package files were changed, which could indicate that a valid binary was overwritten with a trojaned binary.

Perform the following to verify integrity of installed packages

```
rpm -Va | awk '$2 != "c" { print $0}'
```

If any output shows up, you may have an integrity issue with that package that must be rectified.

6.10 VERIFY FILE PERMISSIONS USING RPM

The RPM package management system can check file access permissions of installed software packages, including many that are important to system security. The following command will list which only those files on the system have verification failure due to setting that are different than what are expected by the RPM database:

```
# rpm -Va | grep '^.M'
```

Where:

- `s` is the file size.
- `m` is the file's mode.
- `5` is the MD5 checksum of the file.
- `D` is the file's major and minor numbers.
- `L` is the file's symbolic link contents.
- `U` is owner of the file.
- `G` is the file's group.
- `T` is the modification time of the file.
- `c` appears only if the file is a configuration file. This is handy for quickly identifying config files, as they are very likely to change, and therefore, very *unlikely* to verify successfully.
- `file` is the file that failed verification. The complete path is listed to make it easy to find.

Permissions on system binaries and configuration files that are too generous could allow an unauthorized user to gain privileges that they should not have. The permissions set by the vendor should be maintained. Any deviations from this baseline should be investigated

6.11 ADVANCED INTRUSION DETECTION ENVIRONMENT (AIDE)

AIDE is a file integrity checking tool, similar in nature to Tripwire. While it cannot prevent intrusions, it can detect unauthorized changes to configuration files by alerting when the files are changed. When setting up AIDE, decide internally what the site policy will be concerning integrity checking. Review the AIDE quick start guide and AIDE documentation before proceeding.

6.11.1 Install AIDE

Install AIDE to make use of the file integrity features to monitor critical files for changes that could affect the security of the system. The AIDE package must be installed during the initial kick start of the system.

Perform the following command to verify that AIDE is installed

```
# rpm -q aide
aide-<version>.el7.x86_64
```

6.11.2 Initialize AIDE

To initialize AIDE, run the following command


```
/usr/sbin/aide --init -B 'database_out=file:/var/lib/aide/aide.db.gz'
```

6.11.3 Implement Periodic Execution of File Integrity Check

➤ Associated Baseline Configuration Files

`/etc/cron.d/aide-check`

`/etc/sysconfig/prelink`

Periodic file checking allows the system administrator to determine on a regular basis if critical files have been changed in an unauthorized fashion.

Verify that the default configuration file, `/etc/cron.d/aide-check`, contains the following lines:

```
#/etc/cron.d/aide-check
# modified by:<user name>
0 5 * * * /usr/sbin/aide -check
```

6.11.3.1 Disable prelinking

The prelinking feature changes binaries in an attempt to decrease their startup time. The prelinking feature can interfere with Advanced Intrusion Detection Environment (AIDE) because of the fact that it, prelinking, alters binaries. In order to disable it, change the following line inside the file `/etc/sysconfig/prelink`

```
PRELINKING=no
```

Next, run the following command to return binaries to a normal, non-prelinked state:

```
/usr/sbin/prelink -ua
```

6.12 DISABLE UNUSED SERVICES

Disabling unnecessary services within Red Hat Enterprise Linux greatly reduces the risk of security vulnerabilities.

- Disable the following services
- If any of these services are required for a particular application or purpose, then it must be documented

aep1000	amd	apmd	arptables_jf	atd	autofs	bcm5820
bgpd	canna	chargen	daytime	dc_client	dc_server	dhcpcd
dhcrelay	finger	FreeWnn	gmp	gpm	hpoj	identd
innd	irda	irqbalance	isdn	kadmin	kprop	krb524
krb5kdc	kudzu	lisa	mdmonitor	microcode_ctl	mysqld	named
netdump	netdump-server	netfs	nfs	nfsloc	nfslock	ospf6d
ospfd	pcmcia	portmap	postgresql	psacct	pxe	qotd
radiusd	radvd	rarpd	ripd	ripngd	rlogin	rsh
rstatd	rusersd	rwhod	sendmail	smartd	smb	snmptrapd
spamassassin	squid	telnet	tux	vncserver	vsftpd	winbind
	yum-updatesd	yppasswdd	yppasswd	ypserv	ypxfrd	zebra

- The script below will disable unnecessary services. For any service that are not enabled there will be an error “Failed to issue method call: Access denied “, which is to be expected:

```
for i in echo daytime finger chargen qotd atd gmp kudzu netfs nfsloc portmap
identd netdump netdump-server rwhod sendmail rlogin smb yppasswd ypserv
ypxfrd rsh telnet apmd gpm autofs isdn nfs nfslock winbind irda mysqld
postgresql microcode_ctl mdmonitor psacct pcmcia irqbalance smartd snmptrapd
vncserver hpoj spamassassin lisa canna amd dc_client dc_server aep1000
bcm5820 squid dhcrelay rstatd rusersd FreeWnn named vsftpd pxe rarpd kadmin
kprop krb524 krb5kdc dhcpcd radiusd yppasswdd bgpd ospf6d ospfd ripd ripngd
zebra radvd arptables_jf innd tux yum-updatesd
do
systemctl disable $i
done
```

Additionally, delete all unnecessary xinetd services from the /etc/xinetd.d directory.

This includes services such as:

telnet, ftp, time, gss* krb* ekl* klog* chargen* kshell* daytime* echo* time* cups* sgi* ktalk
Amanda rexec rlogin dbs* aman* amidxtape tftp imap* pop

6.12.1 Disable Interactive startup

To prevent users from starting up the system interactively, as root, disable the PROMPT parameter in the /etc/sysconfig/init file:

```
# PROMPT=no
```

6.12.2 Verify grub.cfg settings

The file `/etc/boot/grub2/grub.cfg` should be owned by the root user and group and permissions set to 600 to prevent destruction or modification of the file. This is the default setting, but to ensure that these settings are correct run the commands:

```
# stat -L -c "%u %g" /boot/grub2/grub.cfg | egrep "0 0"  
0 0
```

Set permission on the `/boot/grub2/grub.cfg` file to read and write for root only

```
# chmod og-rwx /boot/grub2/grub.cfg
```

6.12.3 Shadow Passwords

Implement the use of shadow passwords. This is installed by default on Red Hat Linux

6.12.4 Verify that no Non-Root Accounts have a UID of 0

Only the root user account should have a UID of 0. If any other accounts have a UID of 0, ensure that these additional UID-0 accounts are authorized, and that there is a good reason for them to exist.

The following command will print all password file entries for accounts with UID 0:

```
# awk -F: '($3 == "0") {print}' /etc/passwd
```

6.12.5 Verify settings on /etc/shadow

To properly set the owner, group and settings of `/etc/shadow`:

```
# chown root /etc/shadow  
# chgrp root /etc/shadow  
# chmod 0000 /etc/shadow
```

6.12.6 Verify settings on /etc/gshadow

To properly set the owner, group and settings of `/etc/shadow`:

```
# chown root /etc/gshadow  
# chgrp root /etc/gshadow  
# chmod 0000 /etc/gshadow
```

6.12.7 Verify settings on /etc/group

To properly set the owner, group and settings of /etc/shadow:

```
# chown root /etc/group
# chgrp root /etc/group
# chmod 0000 /etc/group
```

6.12.8 Verify settings on /etc/passwd

To properly set the owner, group and settings of /etc/shadow:

```
# chown root /etc/passwd
# chgrp root /etc/passwd
# chmod 0000 /etc/passwd
```

6.12.9 Ensure Password File Integrity

Ensure there are no mistakes in the /etc/passwd or /etc/groups files:

Verify that there are no accounts with blank passwords.

Run the following commands to ensure there are no mistakes in the /etc/passwd or /etc/groups files:

```
# /usr/sbin/pwck
# /usr/sbin/grpck
```

To verify that there are no accounts with blank passwords run the following command:

```
awk -F: '($2 == "") { print $1 }' /etc/shadow
```

6.12.10 Lock Inactive User Accounts

VA guidelines specify that unused user accounts must be locked after 90 days of inactivity.

Use the passwd command to lock an inactive user account.

To identify user accounts that are over 90 days since they were last logged into, use the command below:

```
lastlog -b 90
```

For user accounts that have the entry ***Never logged in*** for the last login time, use the command below to determine when the account was created and if it is over 90 days old, it must be disabled:

```
passwd -S <username>
```

i Note: It is expected that system accounts and service accounts will have the entry ***Never Logged in*** for the last login time

6.13 ACCOUNT MANAGEMENT

Individual user accountability precludes the use of shared accounts (i.e., accounts where multiple users are allowed to log on directly to the same account). Applications may require that a specific account (e.g., oracle) be used for certain administrative tasks. The user is required to log on with that user's account name and sudo to the application account. That action retains the individual accountability (through audit files). If there is an application account (e.g., oracle) that requires the account to be shared, this will be justified and documented.

6.13.1 Mandatory Password Management Practices

The following specify the mandatory password management practices of all VA Information systems enterprise wide.

6.13.1.1 Password must not include any of the following:

- Vendor/manufacturer default passwords
- Names (i.e. systems user names, part or entire account name of user, family names)
- Words found in the dictionary (i.e. words from any dictionary; spelled forward or backwards)
- Addresses
- Birthdays
- Social Security Numbers
- Common character sequences (i.e. 3456, ghjk, 2468)
- Vendor supplied default passwords (i.e. SYSTEM, Password, Default, USER, Demo and TEST)

6.13.1.2 Passwords must :

- Have at least eight characters
- Administrator accounts shall use a password with a minimum of 12 characters
- contain characters from three of the following categories
 - English upper case characters (A-Z)

- English lower case characters (a-z)
- Base 10 digits (0-9)
- Non-alphanumeric, special characters (i.e. !@#\$%^&*)
- Six of the characters must not occur more than once in the password

6.13.1.3 Passwords must:

- Be changed at least every 90 days
- Be changed immediately, if believed to be compromised or one suspects a password has been compromised, or if discovered to be in non-compliance with this standard
- Be changed on direction from management
- Expire after 90 days of inactivity (System administrator accounts would be exempt from expiring due to inactivity)
- Have minimum of 2 days between password changes
- Users will receive warnings beginning 14 days before their password expires
- Do not reuse a password you have used during the past 5 password changes, or as recent as two years from when the password was last used in order to preclude password guessing
- An intruder lock out feature shall suspend the account after 5 invalid attempts to log on to the system
- Where round-the-lock system administration service is available, system administrator intervention shall be required to clear a locked account
- Where round-the-clock system administration is not available, accounts shall remain locked out for at least fifteen (15) minutes

6.13.2 Service Account Management

The follow applies to all Linux service accounts.

- Service accounts shall use a password with a minimum of 12 characters in length
- A password randomizing tool shall not be used where Denial of Service (DoS) risks exist
- Password shall be changed every 3 years (at Certification and Accreditation (C and A), Security Control Assessment (SCA) test time)
- The *mkpasswd* utility (available from the expect package) can be used to generate and apply password to service accounts.

The following command generates a random 12 character password with at least one digit, one lower case character, one upper case character and one special character and assigns it to the user. Sending output to /dev/null prevents the result from echoing back to the screen

Use the script below to disable service accounts, ensure they have no login shell and that the password aging requirement is set

```
#Lock and Set strong password on service accounts
#The expect package must be installed
servaccts=$( awk -F: '($3 >0 && $3< 500) { print $1 }' /etc/passwd )
for acct in $servaccts
do
mkpasswd -l 12 -d 1 -c 1 -C 1 -s 1 $acct >/dev/null
/usr/sbin/usermod -L -s /sbin/nologin $acct
passwd -x 1095 -n 2 -w 14 -i 30 $acct
done
```

6.13.3 Creating Unique Local User Accounts

VA directives require unique identification for each system user. Authorized users should be granted access only to the resources needed to accomplish their mission. A user is either an individual or an executing process/task that accesses a computer resource. Each user will be identified with an account name and a corresponding user identification (UID) number. The uid's and group identification (gid) numbers are assigned according to the following scheme

Although Linux systems automatically generate unique user ID's (UID) and group ID's (GID) for each user when the account is created on the individual system, it is preferable to ensure that the UID and GID are the same on all systems. The following steps can be used to accomplish this.

The script below is used to create a unique uid based on the user's active directory name. Before creating the user account add a 1 and another digit according to the following legend.

0 = National Accounts

1 = Region 1

2 = Region 2

3 = Region 3

4 = Region 4

5 = Region 5

6= Region 6

Create user accounts that are the same as the users Active Directory account and if they have a 0 account, use that.

In this case below, the ID for a national account would be 109491, for region 6 it would be 169491.

Create a unique 4-digit ID for the user. Substitute the real username for the *VHATESTERJ0* name in the following script. **Note:** that all capital letters are converted to lowercase.

```
name="VHATESTERJ0"; myname=$(echo "${name}" | awk -F^ '{ print tolower($1) }')
num=$((echo -n $myname) | md5sum | cut -d' ' -f1 | tr -d 'a-f' | cut -c1-4); echo
$myname:"$num

>>output = vhatesterj0:9491
```

Create an encrypted password for the account.

i This will create an md5 encrypted password, but once the user changes their password at first logon, it will be sha512 encrypted.

```
openssl passwd -1
Password: <Enter password here>
Verifying - Password:<Re-enter password>
$1$amfJDQQB$8BxcE20oBs5pZ2Jo/kZnI/ <<encrypted password Purple#2012>>
```

Create unique group for the user

```
/usr/sbin/groupadd -g 119491 vhatesterj0
```

Create the user and add them to their own group, a previously created *r01linuxadmins* group and the built-in *wheel* group. The last thing is to force the password to expire so that the user must change it the first time they log in.

i **Note:** The comment field (-c) can be anything, but it is recommended to use the email address of the user to make it clear who that user is.

```
/usr/sbin/useradd -u 119491 \
-c 'john.tester@va.gov' \
-d /home/vhatesterj0 \
-m \
-g vhatesterj0 \
-G r01linuxadmins,wheel \
-p '$1$amfJDQQB$8BxcE20oBs5pZ2Jo/kZnI/' \
-s /bin/bash \
vhatesterj0

#Force the password to expire
chage -d 0 vhatesterj0
```


6.13.4 The SUDO command must require authentication

- Associated Baseline Configuration Files

/etc/sudoers

The "sudo" command allows authorized users to run programs (including shells) as other users, system users, and root. Some configuration options in the "/etc/sudoers" file allow configured users to run programs without re-authenticating. Use of these configuration options makes it easier for one compromised account to be used to compromise other accounts

Edit the */etc/sudoers* file and verify that the following lines are set correctly

```
***** Comment '#' removed from this line
%wheel  ALL=(ALL)          ALL

***** Comment '#' does exist in front of this line
# %wheel  ALL=(ALL)        NOPASSWD: ALL
```

6.14 RHEL CONFIGURATION CHANNELS

A configuration channel is a collection of files that are deployed to the systems at installation time, and can also be redeployed as necessary to keep systems up to date and are deployed by the National Satellite Server. These configuration channels are used to configure the settings that are set forth as requirements in this guide. All Red Hat Linux systems are subscribed to at least the following configurations channels. Each region will replace the VA or va text, with the appropriate region:

- VA baseline RHEL7 (va-baseline-rhel7)
/etc/at.allow
/etc/bashrc
/etc/cron.allow
/etc/cron.d/aide-check
/etc/cron.daily/rdp_update
/etc/csh.login
/etc/inittab
/etc/issue (symlink to /etc/issue.net)
/etc/issue.net

```

/etc/login.defs
/etc/logrotate.conf
/etc/logrotate.d/syslog
/etc/ntp.conf
/etc/ntp/step-tickers
/etc/pam.d/su
/etc/pam.d/system-auth (symlink to /etc/pam.d/system-auth-va)
/etc/pam.d/system-auth-va
/etc/profile.d/tmout.csh
/etc/profile.d/tmout.sh
/etc/securetty
/etc/sysconfig/init
/etc/modprobe.d/crisp.conf
/etc/sysconfig/prelink
/etc/sysctl.conf
/etc/sysctl.d/
/etc/security/limits.d/50-crisp.conf
/etc/selinux/config
/etc/ssh/sshd_config
/etc/sudoers

```

- VA V2S (va-v2s)


```

/etc/opt/BESClient/actionsite.afxm
/etc/snmp/snmpd.conf
/root/EPO_4.X/RX_install.sh ←where X is the current version
/root/RX_install.sh
/root/v2s_kick.sh

```

6.14.1 VA Baseline RHEL 7

6.14.1.1 Cron and Crontab

- Associated Baseline Configuration Files

```

/etc/at.allow
/etc/cron.allow

```

6.14.1.2 Enable cron Service

The cron service is used to execute commands at preconfigured times. It is required by almost all systems to perform necessary maintenance tasks, such as notifying root of system activity. The *crond* service can be enabled with the following command:

```
systemctl enable crond
```

It is required to use files in */etc/cron.d*, */etc/cron.daily*, */etc/cron.hourly*, */etc/cron.monthly*, and */etc/cron.weekly* for scheduling system jobs rather than root's crontab unless there is a compelling reason not to do so and will require a Risk Based Decision (RBD).

6.14.1.3 Deny User Access to CRONTAB

The */etc/at.allow* and */etc/cron.allow* files contain lists of users who are allowed to use cron and at to delay execution of processes. If these files exist and if the corresponding files */etc/cron.deny* and */etc/at.deny* do not exist, then only users listed in the relevant allow files can run the crontab and at commands to submit jobs to be run at scheduled intervals.

On most systems, only the system administrator needs the ability to schedule jobs. **Note:** even if a given user is not listed in *cron.allow*, cron jobs can still be run as that user. The *cron.allow* file controls only administrative access to the *crontab* command for scheduling and modifying cron jobs..

- Edit */etc/cron.allow*, adding one line for each user allowed to use the crontab command to create cron jobs.
- Edit */etc/at.allow*, adding one line for each user allowed to use the “at” command to create jobs.
- Remove the */etc/cron.deny* file
- Remove the */etc/at.deny* file.

6.14.2 Modify Default System-wide Profiles

➤ Associated Baseline Configuration Files

/etc/bashrc

/etc/csh.cshrc

/etc/profile

/etc/login.defs

/etc/csh.login

The umask setting controls the default permissions for the creation of new files. With a default umask setting of 077, files and directories created by users will not be readable by any other user on the system. Users who wish to make specific files group- or world-readable can accomplish this by using the chmod command.

Additionally, users can make all their files readable to their group by default by setting a umask of 027 in their shell configuration files.

6.14.2.1 Ensure Default Umask settings

To ensure the default umask for users of the Bash shell is set properly to 077 in */etc/bashrc*, use the following script:

```
#!/bin/bash
#
user_umask="077"
grep -q umask /etc/bashrc && \
    sed -i "s/umask.*/umask $user_umask/g" /etc/bashrc
if ! [ $? -eq 0 ]; then
    echo "umask $user_umask" >> /etc/bashrc
fi
```

To ensure the default umask for users of the C shell is set properly to 077 in */etc/csh.cshrc* use the following script:

```
#!/bin/bash
user_umask="077"
grep -q umask /etc/csh.cshrc && \
    sed -i "s/umask.*/umask $user_umask/g" /etc/csh.cshrc
if ! [ $? -eq 0 ]; then
    echo "umask $user_umask" >> /etc/csh.cshrc
fi
```

To ensure the default umask is set properly, add or correct the umask setting in */etc/profile* to read as follows:

```
umask 077
```

To ensure the default umask controlled by */etc/login.defs* is set properly, add or correct the UMASK setting in */etc/login.defs* to read as follows:

```
UMASK 077
```

To ensure the default umask controlled by `/etc/csh.login` is set properly, add or correct the `UMASK` setting in `/etc/csh.login` to read as follows:

```
umask 077
```

6.14.3 Disable X Windows startup at boot time

The X Windows system provides a Graphical User Interface (GUI) where users can have multiple windows in which to run programs and various add on. To prevent automatic startup of X windows, run the following command:

```
systemctl set-default multi-user.target
```

6.14.4 Configuring Secure Sockets Layer (SSH) Server

➤ Associated Baseline Configuration Files

```
/etc/ssh/sshd_config  
/etc/issue.net  
/etc/issue
```

SSH provides confidentiality and integrity for data exchanged between two systems, as well as server authentication, through the use of public key cryptography. The implementation included with the system is called OpenSSH.

6.14.4.1 Allow Only SSH Protocol 2

Only SSH protocol version 2, or higher connections are permitted as protocol version 1 suffers from design flaws that result in security vulnerabilities. The default setting in `/etc/ssh/sshd_config` is correct, and can be verified by ensuring that the following line appears:

```
Protocol 2
```

6.14.4.2 Disable SSH Access via Empty Passwords

To explicitly disallow remote login from accounts with empty passwords, verify that the default configuration file `/etc/ssh/sshd_config` contains the following entry:

```
PermitEmptyPasswords no
```

Any accounts with empty passwords should be disabled immediately and PAM configuration should prevent users from being able to assign themselves empty passwords.

6.14.4.3 Do not permit root login via ssh

The root user should never be allowed to login to a system directly over a network. To disable root login via SSH, add or correct the following line in `/etc/ssh/sshd_config`:

```
PermitRootLogin no
```

6.14.4.4 Do Not Allow SSH Environment Options

To ensure users are not able to present environment options to the SSH daemon, add or correct the following line in `/etc/ssh/sshd_config`:

```
PermitUserEnvironment no
```

6.14.4.5 Configure SSH Warning Banner

Verify that the default configuration file `/etc/issue.net`, to contains the following text.

```
Security Warning!
```

```
This U.S. government system is intended to be used by authorized VA network users for
viewing and retrieving
information only except as otherwise explicitly authorized. VA information resides on
and transmits through computer systems and networks funded by the VA. All use is
considered to be with an understanding and acceptance that there is no reasonable
expectation of privacy for any data or transmissions on Government Intranet or
Extranet (non-public) networks or systems.,
All transactions that occur on this system and all data transmitted through this
system are subject to review and action including (but not limited to) monitoring-
recording- retrieving- coping- auditing- inspecting- investigating- restricting
access- blocking- tracking- disclosing to authorized personnel or any other
authorized actions by all authorized VA and law enforcement personnel.,
```

```
All use of this system constitutes understanding and unconditional acceptance of
these terms., Unauthorized attempts or acts to either (1) access- upload- change- or
delete information on this system(2) modify this system (3) deny access to this
system or (4) accrue resources for unauthorized use on this system are strictly
prohibited., Such attempts or acts are subject to action that may result in criminal
civil or administrative penalties
```

Add a symlink from `/etc/issue.net` to `/etc/issue`

Next, verify that the default configuration file `/etc/ssh/sshd_config` contains the following text.

```
Banner /etc/issue.net
```

6.14.4.6 Use Only Approved Ciphers

Limiting the ciphers to those algorithms which are FIPS-approved, Counter (CTR) mode is also preferred over cipher-block chaining (CBC) mode.

The “built-in” default Ciphers are:

```
aes128-ctr,aes192-ctr,aes256-ctr,arcfour256,arcfour128,aes128-cbc,3des-cbc,blowfish-cbc,cast128-cbc,aes192-cbc,aes256-cbc,arcfour
```

The arcfour and blowfish ciphers are not fips certified. Verify that the default configuration file `/etc/ssh/sshd_config` contains the entry as below and if not, it must be added.

```
#Use only fips certified ciphers
Ciphers aes128-ctr,aes192-ctr,aes256-ctr,aes128-cbc,3des-cbc,aes192-cbc,aes256-cbc

#Disable MD5 MACs (Message Authentication Code) Algorithms
MACs hmac-sha1,hmac-ripemd160
```

6.14.5 Configure /etc/login.defs to meet password default requirements

➤ Associated Baseline Configuration Files

`/etc/login.defs`

Verify that the baseline configuration `/etc/login.defs` file has the following entries:.

```
PASS_MAX_DAYS      90
PASS_MIN_DAYS      2
PASS_MIN_LEN        12
PASS_WARN_AGE       14
FAIL_DELAY          15
```

6.14.6 Password Protect Single User Mode

➤ Associated Baseline Configuration Files

VA Baseline Configuration and Security Standard RHEL 7

```
/etc/sysconfig/init
```

Single-user mode is intended as a system recovery method, providing a single user root access to the system by providing a boot option at startup. By default, no authentication is performed if single-user mode is selected.

To require entry of the root password even if the system is started in single-user mode, add or correct the following line in the file `/etc/sysconfig/init`:

```
SINGLE=/sbin/sulogin
```

6.14.7 Logging

➤ Associated Baseline Configuration Files

```
/etc/rsyslog.conf
```

```
/etc/logrotate.d/rsyslog
```

The *rsyslog* software is a replacement for the default *syslogd* daemon and provides the option to log to database formats, and the encryption of log data to a central logging server. *rsyslog* is installed by default on RHEL 7.

6.14.7.1 Configuring logging rules

The `/etc/rsyslog.conf` file specifies rules for logging and which files are to be used to log certain classes of messages.

Review the contents of the `/etc/rsyslog.conf` file to ensure appropriate logging is set.

Custom logging configuration files should be created in the `/etc/rsyslog.d` directory.

6.14.7.2 Log Rotation

Red Hat comes with a *logrotate* program built-in. The log files should be rotated at least monthly. However the frequency of the rotation of those, and other log files, will vary depending upon the logs being kept and the rotation will need to comply with VA/regional log retention regulations.

6.14.7.3 Log all sudo actions

➤ Associated Baseline Configuration Files

```
/etc/rsyslog.d/sudo.conf
```



```
/etc/logrotate.d/sudoers
/etc/sudoers
```

Create the `/etc/rsyslog.d/sudo.conf` file and add the following line:

```
local2.debug /var/log/sudo.log
```

Modify the `/etc/sudoers` file to add the following entry

```
Defaults logfile=/var/log/sudo.log,loglinelen=0
```

Create the file `/etc/logrotate.d/sudoers` and add the following entries:

```
/var/log/sudoe.log {
    weekly
    missingok
    rotate 4
    compress
    delaycompress
    copytruncate
    minsize 100k
}
```

Additionally, view the file `/etc/logrotate.d/syslog` and verify that the names of all of the log files listed in the `/etc/rsyslog.conf` file are present so they will all be rotated.

6.14.8 Disable Unused Kernel Drivers

6.14.8.1 Disable Modprobe Loading of USB Storage Driver

➤ Associated Baseline Configuration Files

```
/etc/modprobe.d/crisp.conf
```

To prevent USB storage devices from being used, configure the kernel module loading system to prevent automatic loading of the USB storage driver.

To configure the system to prevent the `usb-storage` kernel module from being loaded, add the following line to `/etc/modprobe.d/crisp.conf`:

```
install usb-storage /bin/false
```

This will prevent the *modprobe* program from loading the *usb-storage* module, but will not prevent an administrator (or another program) from using the "insmod" program to load the module manually if necessary.

6.14.8.2 Disable Modprobe DCCP Support

- Associated Baseline Configuration Files

`/etc/modprobe.d/crisp.conf`

To configure the system to prevent the Datagram Congestion Control Protocol (DCCP) kernel module from being loaded, add the following line to `/etc/modprobe.d/crisp.conf`:

```
install dccp /bin/false
```

6.14.8.3 Disable Modprobe SCTP Support

- Associated Baseline Configuration Files

`/etc/modprobe.d/crisp.conf`

To configure the system to prevent the Stream Control Transmission Protocol (SCTP) kernel module from being loaded, add the following line to `/etc/modprobe.d/crisp.conf`:

```
install sctp /bin/false
```

6.14.8.4 Disable Modprobe RDS Support

- Associated Baseline Configuration Files

`/etc/modprobe.d/crisp.conf`

To configure the system to prevent the Reliable Datagram Sockets (RDS) kernel module from being loaded, add the following line to `/etc/modprobe.d/crisp.conf` :

```
install rds /bin/false
```

6.14.8.5 Disable Modprobe TIPC Support

- Associated Baseline Configuration Files

`/etc/modprobe.d/crisp.conf`

To configure the system to prevent the Transparent Inter-Process Communication (TIPC) kernel module from being loaded, add the following line to `/etc/modprobe.d/crisp.conf` :

```
install tipc /bin/false
```

6.14.8.6 Disable Modprobe Bluetooth Support

- Associated Baseline Configuration Files

`/etc/modprobe.d/crisp.conf`

To configure the system to prevent the Bluetooth kernel module from being loaded, create and add the following line to `/etc/modprobe.d/crisp.conf` :

```
install bluetooth /bin/false
```

6.14.8.7 Disable cramfs module for loading

- Associated Baseline Configuration Files

`/etc/modprobe.d/crisp.conf`

To configure the system to prevent the cramfs kernel module from being loaded, add the following line to `/etc/modprobe.d/crisp.conf` :

```
install cramfs /bin/false
```

6.14.8.8 Disable freevxfs module for loading

- Associated Baseline Configuration Files

`/etc/modprobe.d/crisp.conf`

To configure the system to prevent the freevxfs kernel module from being loaded, add the following line to `/etc/modprobe.d/crisp.conf` :

```
install freevxfs /bin/false
```

6.14.8.9 Disable jffs2 module for loading

➤ Associated Baseline Configuration Files

`/etc/modprobe.d/crisp.conf`

To configure the system to prevent the jffs2 kernel module from being loaded, add the following line to `/etc/modprobe.d/crisp.conf` :

```
install jffs2 /bin/false
```

6.14.8.10 Disable hfs module for loading

➤ Associated Baseline Configuration Files

`/etc/modprobe.d/crisp.conf`

To configure the system to prevent the hfs kernel module from being loaded, add the following line to `/etc/modprobe.d/crisp.conf` :

```
install hfs /bin/false
```

6.14.8.11 Disable hfsplus module for loading

➤ Associated Baseline Configuration Files

`/etc/modprobe.d/crisp.conf`

To configure the system to prevent the hfsplus kernel module from being loaded, add the following line to `/etc/modprobe.d/crisp.conf` :

```
install hfsplus /bin/false
```

6.14.8.12 Disable squashfs module for loading

➤ Associated Baseline Configuration Files

`/etc/modprobe.d/crisp.conf`

To configure the system to prevent the squashfs kernel module from being loaded, add the following line to `/etc/modprobe.d/crisp.conf` :

```
install squashfs /bin/false
```

6.14.8.13 Disable udf module for loading

➤ Associated Baseline Configuration Files

`/etc/modprobe.d/crisp.conf`

To configure the system to prevent the udf kernel module from being loaded, add the following line to `/etc/modprobe.d/crisp.conf` :

```
install udf /bin/false
```

6.14.9 System Time Protocol Service (NTP)

➤ Associated Baseline Configuration Files

`/etc/ntp.conf`

`/etc/ntp/step-tickers`

Network Time Protocol is used to manage the system clock over a network. Time services are very important for logs as well as such things as Secure Socket Layer (SSL) certificates, so it is essential that the time be synchronized within an environment.

6.14.9.1 Specify Network Time Protocol (NTP) Server

To specify a remote NTP server for time synchronization, verify that the following lines exist in the default configuration file `/etc/ntp.conf`. Although the settings look the same, this will ensure that 3 different time servers are configured

```
server ntp.va.gov -iburst
server ntp.va.gov -iburst
server ntp.va.gov -iburst
```

6.14.9.2 Enable and start the NTP Daemon

The ntpd service is enabled and started with the following commands:

```
systemctl enable ntpd.service
systemctl start ntpd.service
```

To sync the system with the NTP server run the following command.

```
ntpdate -gq
```

6.14.10 Configuring PAM for user account compliance

➤ Associated Baseline Configuration Files

```
/etc/pam.d/system-auth-ac
/etc/pam.d/system-auth-va
/etc/pam.d/password-auth-ac
/etc/pam.d/passwd-auth-va
/etc/pam.d/system-auth
/etc/security/pwquality.conf
```

Pluggable Authentication Modules (PAM) is a system which implements modular authentication for Linux programs. PAM provides a flexible and configurable architecture for authentication, and it should be configured to minimize exposure to unnecessary risk. This section contains guidance on how to accomplish that.

The Authentication Configuration Tool automatically writes to the `/etc/pam.d/system-auth-ac` file, which is symlinked to `/etc/pam.d/system-auth`. Any changes made to `/etc/pam.d/system-auth` are overwritten the next time that `authconfig` is ran.

The following commands will remove the symlink for `/etc/pam.d/system-auth`, copy the `/etc/pam.d/system-auth-ac` file to `/etc/pam.d/system-auth-va`, and create a symlink from `/etc/pam.d/system-auth-va` to `/etc/pam.d/system-auth`.

And also remove the symlink for `/etc/pam.d/password-auth`, copy the `/etc/pam.d/password-auth-ac` file to `/etc/pam.d/password-auth-va`, and create a symlink from `/etc/pam.d/password-auth-va` to `/etc/pam.d/password-auth`.

```
rm /etc/pam.d/system-auth
#
cp /etc/pam.d/system-auth-ac /etc/pam.d/system-auth-va
#
ln -s /etc/pam.d/system-auth-va /etc/pam.d/system-auth
```

```
#
rm /etc/pam.d/password-auth
#
cp /etc/pam.d/password-auth-ac /etc/pam.d/password-auth-va
#
ln -s /etc/pam.d/password-auth-va /etc/pam.d/password-auth
```

6.14.10.1 PAM settings For password complexity using pam_cracklib

The *pam_cracklib* module checks of the strength of passwords. It performs checks such as making sure a password is not a dictionary word, it is a certain length, contains a mix of characters (e.g. alphabet, numeric, other) and more.

The PAM module config file, */etc/security/pwquality.conf*, will be configured with the settings below to ensure password complexity.

- Set password Minimum length minlen = 12
- Set Password to Maximum of
 - Allowed Consecutive Repeating Characters maxrepeat = 3
 - Allowed consecutive characters of same class maxclassrepeat = 3
- Set Password Strength Minimum Digit Characters dcredit = -1
- Password Strength Minimum Uppercase Characters ucredit = -1
- Set Password Strength Minimum Special Characters ocredit = -1
- Set Password Strength Minimum Lowercase Characters lcredit = -1
- Set Password Strength Minimum Different Characters difok = 5

6.14.10.2 PAM settings for failed password attempts

The PAM module config file, */etc/pam.d/system-auth-va* and */etc/pam.d/password-auth-va* will be configured with the settings below for failed password attempts. The *pam_faillock* PAM module provides the capability to lock out user accounts after a number of failed login attempts

- Set deny for failed password attempts
 - Deny=5
 - Unlock_time=900 (15 minutes)
 - Fail_interval=900 (15 minutes)

Add the following lines immediately **before** the *pam_env.so* statement in */etc/pam.d/system-auth-va* and */etc/pam.d/password-auth-va*:

```
auth required pam_faillock.so preauth silent deny=5 unlock_time=900 fail_interval=900
```

Add the following lines immediately **after** the `pam_env.so` statement in `/etc/pam.d/system-auth-vb` and `/etc/pam.d/password-auth`

```
auth [default=die] pam_faillock.so authfail deny=5 unlock_time=900 fail_interval=900
```

6.14.10.3 PAM settings to limit password reuse.

PAM can be configured to not allow users to reuse recent passwords. This can be accomplished by using the `remember` option for the `pam_unix` PAM module. In the file `/etc/pam.d/system-auth-vb`, append `remember=5` as below.

```
password sufficient pam_unix.so existing_options remember=5
```

6.14.10.4 PAM settings to display failed logon.

To configure the system to notify users of last logon/access using `pam_lastlog`, add the following line immediately after `session required pam_limits.so` in `/etc/pam.d/system-auth-vb`

```
session      required      pam_lastlog.so showfailed
```

6.14.11 Restrict Root Logins to the System Console

➤ Associated Baseline Configuration Files

`/etc/securetty`

Direct root logins should be allowed only for emergency use. In normal situations, the administrator should access the system via a unique unprivileged account, and use `su` or `sudo` to execute privileged commands. Discouraging administrators from accessing the root account directly ensures an audit trail in organizations with multiple administrators. Locking down the channels through which root can connect directly reduces opportunities for password-guessing against the root account.

The login program uses the file `/etc/securetty` to determine which interfaces should allow root logins.

The virtual devices `/dev/console` and `/dev/tty*` represent the system consoles (accessible via the `Ctrl-Alt-F1` through `Ctrl-Alt-F6` keyboard sequences on a default installation).

Ensure that the `/etc/securetty` file contains only the following lines:

```
#The primary system console device:
```



```
console
#The virtual console devices:
tty1
tty2
tty3
tty4
tty5
tty6
...
```

6.14.12 Restricting use of the SU commands

➤ Associated Baseline Configuration Files

`/etc/pam.d/su`

The `su` command allows a user to run a command or shell as another user. The program has been superseded by `sudo`, which allows for more granular control over privileged access. Normally, the `su` command can be executed by any user. By uncommenting the `pam_wheel.so` statement in `/etc/pam.d/su`, the `su` command will only allow users in the wheel group to execute `su`.

Verify that the default configuration file, `/etc/pam.d/su` contains the line as below.

<code>auth</code>	<code>required</code>	<code>pam_wheel.so use_uid</code>
-------------------	-----------------------	-----------------------------------

Once this is done, use the `usermod` command to add current administrators to the wheel group and future administrators users can be added to that group when their account is first created.

6.14.13 Ensure Password Hashing Algorithm is SHA-512

➤ Associated Baseline Configuration Files

`/etc/pam.d/system-auth-vb`

`/etc/login.defs`

`/etc/libuser.conf`

Configure the system to use the SHA-512 algorithm in the so that new users will have their password hash generated with the SHA-512 algorithm, and when existing users changes their passwords, hashes for the new passwords will be generated using the SHA-512 algorithm. Using a stronger hashing algorithm makes password cracking attacks more difficult.

In order to configure the system to use the SHA-512 algorithm, issue the command below:

<code># /usr/sbin/authconfig --passalgo=sha512 --update</code>
--

The password section of the file `/etc/pam.d/system-auth-vb` controls which PAM modules execute during a password change. Verify that the `pam_unix.so` module in the password section to include the argument `sha512`, as shown below:

```
password    sufficient  pam_unix.so sha512 shadow nullok try_first_pass use_authtok
```

This will help ensure when local users change their passwords, hashes for the new passwords will be generated using the SHA-512 algorithm. This is the default in Red Hat 6.

Verify that the default configuration file, `/etc/login.defs`, contains the line below to ensure the system will use SHA-512 as the hashing algorithm:

```
ENCRYPT_METHOD SHA512
```

Add or correct the `crypt_style` line in the [defaults] section of `/etc/libuser.conf` to ensure the system will use the SHA-512 algorithm for password hashing

```
crypt_style = sha512
```

Using a stronger hashing algorithm makes password cracking attacks more difficult.

6.14.14 Implement Inactivity Timeout for Login Shells

➤ Associated Baseline Configuration Files

```
/etc/profile.d/tmout.sh  
/etc/profile.d/tmout.csh
```

Login shells must be configured to automatically log users out after a period of inactivity.

The following instructions implement a 15-minute idle time-out for the default `bash` and `csh` shells, both files are deployed as part of the default configuration files.

Verify that the default configuration file, `/etc/profile.d/tmout.sh` contains the following lines.

(Using the `EXPORT` option prevents users from modifying their inactivity timeout setting)

```
TMOUT=900  
readonly TMOUT  
export TMOUT
```

Verify that the default configuration file, `/etc/profile.d/tmout.csh` contains the following line.

```
set -r autologout 15
```

6.14.15 Disable ipv6

➤ Associated Baseline Configuration Files

`/etc/sysctl.d/ipv6.conf`

If *IPv6* is not to be used, it is disabled to reduce the attack surface of the system.

To configure the system to prevent the *ipv6* kernel module from being loaded, add the following line to `/etc/sysctl.d/ipv6.conf`:

```
net.ipv6.conf.all.disable_ipv6 = 1
net.ipv6.conf.default.disable_ipv6 = 1
#Disable accepting IPv6 router advertisements
net.ipv6.conf.default.accept_ra = 0
# Disable Accepting IPv6 Redirects
```

In addition, add the following two lines to `/etc/sysconfig/network`.

Note: This file is **NOT** deployed as part of the baseline configuration files as it is unique to each system.

```
NETWORKING_IPV6=no
IPV6INIT=no
```

6.14.16 System Resource Management

➤ Associated Baseline Configuration Files

`/etc/security/limits.d/50-crisp.conf`

These settings are for a baseline only. Systems that host an application, such as Oracle database, Oracle WebLogic, InterSystems Cache, etc, will need to have these setting modified as appropriate and documented in the applicable baseline guide.

6.14.16.1 Limit individual file sizes

Individual file sizes are limited to 100 MB, and a user can only have 150 concurrent processes running. Verify that the default configuration file `/etc/security/limits.d/50-crisp.conf` contains the following entries:

```
# CRISP Baseline
*          hard    core      0
*          hard    fsize     102400
*          hard    nproc     150
root       -       fsize     -1
root       -       nproc     -1
```

These settings are the baseline settings for all systems, but they will require modification depending on the purpose of the system.

For example, if the system is going to be used as a WebLogic Middleware server, settings for the weblogic *user* and *group*, would be configured in a file that would be part of baseline configuration for that system. Create a file `/etc/security/limits.d/60-weblogic.conf` and enter the example settings below. The number 60 in the file name ensures that the settings in this file will override those in the `/etc/security/limits.d/50-crisp.conf` file.

```
@weblogic  hard    nofile     65536
@weblogic  soft    nofile     4096
weblogic   hard    nofile     65536
weblogic   soft    nofile     4096
```

6.14.16.2 Disable core dumps for all users

To disable core dumps for all users, add the following line to `/etc/security/limits.d/50-crisp.conf`:

```
*          hard    core      0
```

To disable core dumps for *SUID* programs, verify that the following line exists in the default configuration file: `/etc/sysctl.d/50-crisp.conf`.

```
#This is the system default.
fs.suid_dumpable = 0
```

Core – A core file can be generated from the memory dump of an executable program. It is generally used to determine why a program aborted. It could also be used to glean confidential information from a core file or unnecessarily occupy large amounts of disk space

Soft set at what is expected to be the maximum number of processes that a user or application should use. If the limit is passed, the system administrator will receive a warning.

Hard: The hard limit is the maximum value that can be reached before the user gets the error messages *Out of file handles*.

6.14.17 Enable Randomized Layout of Virtual Address Space

Address Space Layout Randomization (ASLR) makes it more difficult for an attacker to predict the location of attack code they have introduced into a process's address space during an attempt at exploitation. Additionally, ASLR makes it more difficult for an attacker to know the location of existing code in order to re-purpose it using Return Oriented Programming (ROP) techniques.

Verify that the following line exists in the default configuration file: `/etc/sysctl.d/50-crisp.conf`.

```
kernel.randomize_va_space = 2
```

6.14.18 Configure ExecShield

Execshield is made up of a number of kernel features to provide protection against buffer overflow attacks. These features include prevention of execution in memory data space, and special handling of text buffers.

Verify that the following line exists in the default configuration file `/etc/sysctl.d/50-crisp.conf`.

```
kernel.exec-shield = 1
```

6.14.18.1 Restrict Access to Kernel Message Buffer

To set the runtime status of the `kernel.dmesg_restrict` kernel parameter, add the following line to `/etc/security/limits.d/50-crisp.conf`:

```
kernel.dmesg_restrict = 1
```

6.14.19 Disable Ctrl-Alt-Delete Reboot Activation

A logged-in user who presses *Ctrl-Alt-Del*, when at the console, can reboot the system. If accidentally pressed, as could happen in the case of mixed OS environment, this can create the risk of short-term loss of availability of systems due to unintentional reboot

Enter the following command to disable the *Ctrl-Alt-Del* keyboard sequence to reboot the system

```
systemctl mask ctrl-alt-del.target
```

6.14.20 Preventing Network Attacks

➤ Associated Baseline Configuration Files

`/etc/sysconfig/network`

`/etc/sysctl.d/50-crisp.conf`

6.14.20.1 Disable Zeroconf Networking

Zeroconf networking allows the system to assign itself an IP address and engage in IP communication without a statically-assigned address or even a Dynamic Host Control Protocol (DHCP) server. Automatic address assignment via Zeroconf (or DHCP) is not recommended. To disable Zeroconf automatic route assignment in the 169.254.0.0 subnet, add or correct the following line in `/etc/sysconfig/network`:

```
NOZEROCONF=yes
```

6.14.20.2 Kernel Parameters Which Affect Networking

The `sysctl` utility is used to set parameters which affect the operation of the Linux kernel. Kernel parameters which affect networking and have security implications are described here. Verify that the file `/etc/sysctl.d/50-crisp.conf` contains the following entries:

Disable Kernel Parameter for IP Forwarding,

```
net.ipv4.ip_forward = 0
```

Enable Kernel Parameter to Use TCP Syncookies

```
net.ipv4.tcp_syncookies = 1
```

Enable Kernel Parameter to Use Reverse Path Filtering for All Interfaces

```
net.ipv4.conf.all.rp_filter = 1
```

Enable Kernel Parameter to Use strict mode for Reverse Path Filtering by Default

```
net.ipv4.conf.all.rp_filter = 1  
net.ipv4.conf.all.default_filter = 1
```

Disable Kernel Parameter for Accepting Source-Routed Packets for All Interfaces.

```
net.ipv4.conf.all.accept_source_route = 0
```

Disable Kernel Parameter for Accepting ICMP Redirects for All Interfaces.

```
net.ipv4.conf.all.accept_redirects = 0
```

Disable Kernel Parameter for Accepting ICMP Redirects By Default

```
net.ipv4.conf.default.accept_redirects = 0
```

Enable Kernel Parameter to Ignore ICMP Broadcast Echo Requests

```
net.ipv4.icmp_echo_ignore_broadcasts = 1
```

Enable Kernel Parameter to Ignore Bogus ICMP Error Responses

```
net.ipv4.icmp_ignore_bogus_error_responses = 1
```

Disable Kernel Parameter for Accepting Secure Redirects for All Interfaces

```
net.ipv4.conf.all.secure_redirects = 0  
net.ipv4.conf.all.send_redirects = 0
```

Disable Kernel Parameter for Accepting Source-Routed Packets By Default.

```
net.ipv4.conf.default.accept_source_route = 0
```

Disable Accepting IPv6 Redirects

```
net.ipv6.conf.default.accept_ra = 0
```

Disable Accepting IPv6 Router Advertisements

```
net.ipv6.conf.default.accept_redirects = 0
```

6.14.21 Configure System Auditing with Auditd

➤ Associated Baseline Configuration Files

/etc/audit/auditd.conf

/etc/audit/audit.rules

The audit service provides substantial capabilities for recording system activities. By default, the service audits certain types of security-relevant events such as system logins, account modifications, and authentication events performed by programs such as sudo. Under its default configuration, auditd has modest disk space requirements, and should not noticeably impact system performance.

6.14.21.1 Enable Auditing for services that start at boot

To ensure all processes can be audited, even those which start prior to the audit daemon, add the argument "audit=1" to the kernel line in "/etc/default/grub", in the manner below:

```
GRUB_CMDLINE_LINUX="rd.lvm.lv=rhel/swap crashkernel=auto rd.lvm.lv=rhel/root  
ipv6.disable=1 audit=1"
```

Run the grub2-mkconfig command to regenerate the grub.cfg file.

```
grub2-mkconfig -o /boot/grub2/grub.cfg
```


6.14.21.2 Configure auditd Data Retention

By default, *auditd* rotates 5 logs by size (6MB), retaining a maximum of 30MB of data in total, and refuses to write entries when the disk is too full. This minimizes the risk of audit data filling its partition and impacting other services. There are a number of other parameters that must be configured in */etc/audit/auditd.conf*, other settings should be left as they are.

```
#Number of Logs retained
num_logs = 5
#
#Audit Log size (mb)
max_log_file = 6

#Action upon reaching maximum log size
max_log_action = ROTATE

#Action on Low Disk space
admin_space_left_action = email

#Mail Action on low disk space
action_mail_acct = root
```

6.14.21.3 Configure audit rules

6.14.21.3.1 Record Events That Modify Date and Time Information

Capture events where the system date and/or time has been modified. The parameters in this section are set to determine if the *adjtimex* (tune kernel clock), *settimeofday* (Set time, using *timeval* and *timezone* structures) *stime* (using seconds since 1/1/1970) or *clock_settime* (allows for the setting of several internal clocks and timers) system calls have been executed and always write an audit record to the */var/log/audit.log* file upon exit, tagging the records with the identifier *time-change*

Perform the following to determine if events where the system date and/or time has been modified are captured:

```
# grep time-change /etc/audit/audit.rules
-a always,exit -F arch=b64 -S adjtimex,settimeofday -F key=time-change
-a always,exit -F arch=b64 -S clock_settime -F a0=0x0 -F key=time-change
-w /etc/localtime -p wa -k time-change
```

6.14.21.3.2 Record Events That Modify User/Group Information

Record events affecting the group, *passwd* (user IDs), *shadow* and *gshadow* (passwords) or */etc/security/opasswd* (old passwords, based on remember parameter in the PAM configuration) files. The parameters in this section will watch the files to see if they have been opened for write or have had attribute changes (e.g. permissions) and tag them with the identifier *identity* in the *audit log* file.

```
# grep identity /etc/audit/audit.rules
-w /etc/group -p wa -k identity
-w /etc/passwd -p wa -k identity
-w /etc/gshadow -p wa -k identity
-w /etc/shadow -p wa -k identity
-w /etc/security/opasswd -p wa -k identity
```

6.14.21.3.3 Record Events That Modify the System's Network Environment

Record changes to network environment files or system calls. The below parameters monitor the *sethostname* (set the systems host name) or *setdomainname* (set the systems domainname) system calls, and write an audit event on system call exit. The other parameters monitor the */etc/issue* and */etc/issue.net* files (messages displayed pre-login), */etc/hosts* (file containing host names and associated IP addresses) and */etc/sysconfig/network* (directory containing network interface scripts)

```
# grep system-locale /etc/audit/audit.rules
-a always,exit -F arch=b64 -S sethostname,setdomainname -F key=system-locale
-w /etc/issue -p wa -k system-locale
-w /etc/issue.net -p wa -k system-locale
-w /etc/hosts -p wa -k system-locale
-w /etc/sysconfig/network -p wa -k system-locale
```

6.14.21.3.4 Collect Discretionary Access Control Permission Modification Events

Monitor changes to file permissions, attributes, ownership and group. The parameters in this section track changes for system calls that affect file permissions and attributes. The *chmod*, *fchmod* and *fchmodat* system calls affect the permissions associated with a file. The *chown*, *fchown*, *fchownat* and *lchown* system calls affect owner and group attributes on a file. The *setxattr*, *lsetxattr*, *fsetxattr* (set extended file attributes) and *removexattr*, *lremovexattr*, *fremovexattr* (remove extended file attributes) control extended file attributes. In all cases, an audit record will only be written for non-system userids (*auid* >= 500) and will ignore Daemon events (*auid* = 4294967295). All audit records will be tagged with the identifier *perm_mod*.

Perform the following command and ensure the output is as shown to determine if permission modifications are being recorded:

```
# grep perm_mod /etc/audit/audit.rules
-a always,exit -F arch=b64 -S chmod,fchmod,fchmodat -F auid>=1000 -F auid!=4294967295
-F key=perm_mod
-a always,exit -F arch=b64 -S chown,fchown,lchown,fchownat -F auid>=1000 -F
auid!=4294967295 -F key=perm_mod
-a always,exit -F arch=b64 -S
setxattr,lsetxattr,fsetxattr,removexattr,lremovexattr,fremovexattr -F auid>=1000 -F
auid!=4294967295 -F key=perm_mod
```

6.14.21.3.5 Collect Unsuccessful Unauthorized Access Attempts to Files

Monitor for unsuccessful attempts to access files. The parameters below are associated with system calls that control creation (creat), opening (open, openat) and truncation (truncate, ftruncate) of files. An audit log record will only be written if the user is a non-privileged user (auid >= 500), is not a Daemon event (auid=4294967295) and if the system call returned *EACCES* (permission denied to the file) or *EPERM* (some other permanent error associated

Perform the following command and ensure the output is as shown to determine if there are unsuccessful attempts to access files:

```
# grep access /etc/audit/audit.rules
-a always,exit -F arch=b32 -S open,creat,truncate,openat,open_by_handle_at -F exit=-
EACCES -F auid>=1000 -F auid!=4294967295 -F key=access
-a always,exit -F arch=b32 -S open,creat,truncate,openat,open_by_handle_at -F exit=-
EPERM -F auid>=1000 -F auid!=4294967295 -F key=access
-a always,exit -F arch=b64 -S open,truncate,creat,openat,open_by_handle_at -F exit=-
EACCES -F auid>=1000 -F auid!=4294967295 -F key=access
-a always,exit -F arch=b64 -S open,truncate,creat,openat,open_by_handle_at -F exit=-
EPERM -F auid>=1000 -F auid!=4294967295 -F key=access
```

6.14.21.3.6 Collect Changes to System Administration Scope (sudoers)

Monitor scope changes for system administrations. If the system has been properly configured to force system administrators to log in as themselves first and then use the sudo command to execute privileged commands, it is possible to monitor changes in scope. The file */etc/sudoers* will be written to when the file or its attributes have changed. The audit records will be tagged with the identifier *scope*.

Perform the following to determine if changes to */etc/sudoers* are recorded:

```
# grep scope /etc/audit/audit.rules
-w /etc/sudoers -p wa -k actions
-w /etc/sudoers.d/ -p wa -k actions
```

6.14.22 Configure SELINUX

➤ Associated Baseline Configuration Files

`/etc/selinux/config`

Selinux provides a Mandatory Access Control (MAC) system that greatly augments the default Discretionary Access Control (DAC) model. Under Selinux, every process and every object (files, sockets, pipes) on the system is assigned a security context, a label that includes detailed type information about the object. The kernel allows processes to access objects only if that access is explicitly allowed by the policy in effect. The policy defines transitions, so that a user can be allowed to run software, but the software can run under a different context than the user's default. This automatically limits the damage that the software can do to files accessible by the calling user. The user does not need to take any action to gain this benefit. For an action to occur, both the traditional DAC permissions must be satisfied as well as the Selinux MAC rules. The action will not be allowed if either one of these models does not permit the action. In this way, Selinux rules can only make a system's permissions more restrictive and secure. Selinux requires a complex policy to allow all the actions required of a system under normal operation.

Selinux must be enabled at installation time, but set for *PERMISSIVE* mode.

The Selinux *Permissive* Mode is a state where Selinux permits violation of Selinux policy system wide. In this system wide permissive state policy violations are merely logged.

Verify that the `/etc/selinux/config` file has the following setting

```
SELINUX=permissive
SELINUXTYPE=targeted
```

6.14.23 Mail must be forwarded to one or more system administrators

➤ Associated Baseline Configuration Files

`/etc/aliases`

Some system services utilize email messages sent to the root user to notify system administrators of active or impending issues. These messages must be forwarded to at least one monitored email address. Edit the `/etc/aliases` file to add one or more system administrators email address or a group email address. Once this is done, the command `newaliases` must be run to activate

```
# Person who should get root's mail
#root:      marc
root:       john.aministrator@va.gov
root:       someadministraorgroup@va.gov
```

6.15 VA V2S RHEL7

6.15.1 Configure SNMP

➤ Associated Baseline Configuration Files

/etc/snmp/snmpd.conf

VA Visibility To Servers requires that the SNMP service must be installed for monitoring, best practices should be followed to minimize the security risk from the installation.

- Use SNMP version 2c or greater to comply with VA V2S standards. Earlier versions of SNMP are considered insecure, as they potentially allow unauthorized access to detailed system management information
- Write access to the MIB (Management Information Base) should be allowed only if necessary
- All access to the MIB should be restricted following a principle of least privilege
- Network access should be limited to the maximum extent possible including restricting to expected network addresses both in the configuration files and in the system firewall rules
- Ensure SNMP agents send traps only to, and accept SNMP queries only from, authorized management stations
- Ensure that permissions on the snmpd.conf configuration file (by default, in /etc/snmp) are 640 or more restrictive
- Ensure that any MIB files permissions are also 640 or more restrictive

/etc/snmp/snmpd.conf file format

```
#####
# /etc/snmp/snmpd.conf:
# Last modified by:
#
#####

#### In this section we define what can connect and the community name ####
#          sec.name          source          community
com2sec    VA_Network        10.0.0.0/8      V2S_server_r#
com2sec    local             localhost      V2S_server_r#

### This section defines the security model and connection from above ###
#### Also note we are allowing v2c note
#          groupName         securityModel  securityName
group     ReadOnlyGroup      v2c             local
```

```

group      ReadOnlyGroup    v2c                VA_Network

#### This section is setup to allow what OID can be included or excluded ####
#          name      incl/excl      subtree      mask(optional)
view      all        included        .1          80

# group      context sec.model sec.level prefix read      write      notif
access      ReadOnlyGroup  ""          any      noauth      exact    all      none      none

#### Add your own /etc/snmp/snmpd.local.conf with the following lines ####
#### populated with contact and location information ####
#### Note: snmpd.local.conf will override snmpd.conf setting ####
syslocation      System Location
syscontact        System Email Contact

#####
# SECTION: Agent Operating Mode for InterSystems Cache
# This section defines how the agent will operate when it
# is running.

```

➤ Additional V2S Baseline Configuration Files

```

/etc/opt/BESClient/actionsite.afxm
/root/EPO_4.5/RX_install.sh
/root/RX_install.sh
/root/v2s_kick.sh

```

7 REFERENCES

- National Institute of Standards and Technologies (NIST) *Guide to General Server Security* (SP 800-123) <http://csrc.nist.gov/publications/nistpubs/800-123/SP800-123.pdf>
- National Institute of Standards and Technologies (NIST) *Recommended Security Controls for Federal Information Systems and Organizations* (SP 800-53) <http://csrc.nist.gov/publications/PubsSPs.html#800-53>
- National Institute of Standards and Technologies (NIST) *Guide for Security-Focused Configuration Management of Information Systems* (SP 800 128)
- VA Handbook 6500, *Information Security Program* http://www1.va.gov/vapubs/viewPublication.asp?Pub_ID=56
- VA Memo 120229-005-Baseline Configuration
- ~~Defense Information Systems Administration (DISA) *Security Technical Implementation Guides* (STIG's)~~ *No DISA STIGS have been published as yet.*

8 REVIEW DATE

September 1 2016

9 FOLLOW-UP RESPONSIBILITY

:

APPENDIX A FUNCTIONAL CONFIGURATION BASELINES

INTEL CPU'S

Issue - C-State limits

Recent Linux kernels may have a built-in driver ('intel_idle') which will ignore any C-State limits imposed by Basic Input/Output System (BIOS) or Unified Extensible Firmware Interface (UEFI).

This driver was added to take advantage of the power savings given by C-States on newer Intel Central Processing Units (CPUs).

On systems where latency is an issue, this driver may cause issues by enabling C-States even though they are disabled in the BIOS or UEFI. This can cause minor latency (a few Central Processing Unit (CPU) cycles/nano-seconds) as the CPUs transition out of a C-State and into a running state.

Affected configurations

The system is configured with at least one of the following:

- Red Hat Enterprise Linux 6 update 1, update 2, update 3, update 4
- SUSE Linux Enterprise Server 11 Service Pack 1, Service Pack 2, Service Pack 2 x86_64

This is not vendor specific

This does not affect earlier versions of Red Hat Linux.

Resolution

To disable the 'intel_idle' driver or limit the C-States available to the Operating System , add the following start parameter to the grub kernel

intel_idle.max_cstate=0

DELL SYSTEMS

Issue - Biosdevname

Modern x86 - based Servers support an increasing number of network interface ports on the motherboard in addition to add-in network adapters.

Linux based OSes name these interfaces as ethN. The naming of network interfaces is currently non-deterministic and not governed by any standard in terms of their relationship to the way the ports are wired on the system.

Common user expectations such as 'eth0' representing the first network port on the motherboard as labeled on the server chassis cannot be fulfilled in many cases.

Ensuring that the Ethernet interface names follow the order of the devices as intended by the system designer might not be sufficient. The “ethN” names currently in use do not suggest the Ethernet interface’s physical location, whether it is on the system’s motherboard or if it is on an add in card; or if it is on an add-in card with multiple ports, which port on the card it is on.

Consequently, a naming mechanism that can impart meaning to the network interface’s name based on the physical location of a net

work port in concordance to the intended system design is necessary. To achieve this , the system firmware has the ability to communicate the intended order for network devices on the motherboard to the Operating System via standard mechanisms such as SMBIOS and ACPI.

The new naming scheme, using the 'biosdevname' udev helper utility developed by Dell and released under GPL, suggests new names based on the location of the network adapters on the system as suggested by system BIOS.

The new naming scheme

LAN-On-Motherboard interfaces

em<port number>_< virtual function instance / NPAR Index>

(ethernet-on-motherboard <1,2 ..>)

PCI add-in interfaces

p<slot number>p<port number>_<virtual function instance / NPAR Index>

Affected configurations

The system is configured with at least one of the following:

- Red Hat Enterprise Linux 6 update 1, update 2, update 3, update 4

This is vendor specific.

Dell Systems incorporate this convention by default

Resolution

To disable the this naming convention at installation, the parameter, `biosdevname=0` must be added to the command line parameters for Dell systems.

It can also be added to the command line parameters of other systems without adverse effect.

VIRTUAL MACHINES

Hypervisor specific client tools

Install the client tools that are specific to the hypervisor being used..

APPENDIX B: Default Configuration files

File	/etc/at.allow
Settings	Owner: root Group: root Mode: 644
Contents	root

File	/etc/at.aliases
Settings	Owner: root Group: root Mode: 644
Contents	<pre># # Aliases in this file will NOT be expanded in the header from # Mail, but WILL be visible over networks or from /bin/mail. # # >>>>>>>>>> The program "newaliases" must be run after # >> NOTE >> this file is updated for any changes to # >>>>>>>>>> show through to sendmail. # # Basic system aliases -- these MUST be present. mailer-daemon: postmaster postmaster: root # General redirections for pseudo accounts. bin: root daemon: root adm: root lp: root sync: root shutdown: root halt: root mail: root news: root uucp: root operator: root games: root gopher: root ftp: root nobody: root radiusd: root nut: root dbus: root vcsa: root canna: root wnn: root</pre>

	rpm:	root
	nscd:	root
	pcap:	root
	apache:	root
	webalizer:	root
	dovecot:	root
	fax:	root
	quagga:	root
	radvd:	root
	pvm:	root
	amandabackup:	root
	privoxy:	root
	ident:	root
	named:	root
	xfs:	root
	gdm:	root
	mailnull:	root
	postgres:	root
	sshd:	root
	smmsp:	root
	postfix:	root
	netdump:	root
	ldap:	root
	squid:	root
	ntp:	root
	mysql:	root
	desktop:	root
	rpcuser:	root
	rpc:	root
	nfsnobody:	root
	ingres:	root
	system:	root
	toor:	root
	manager:	root
	dumper:	root
	abuse:	root
	newsadm:	news
	newsadmin:	news
	usenet:	news
	ftpadm:	ftp
	ftpadmin:	ftp
	ftp-adm:	ftp
	ftp-admin:	ftp
	www:	webmaster
	webmaster:	root
	noc:	root
	security:	root
	hostmaster:	root
	info:	postmaster
	marketing:	postmaster
	sales:	postmaster
	support:	postmaster
	# trap decode to catch security attacks	

	<pre> decode: root # Person who should get root's mail root: john.administrator@va.gov root: someadministratorgroup@va.gov </pre>
--	---

File	/etc/bashrc
Settings	<pre> Owner: root Group: root Mode: 644 </pre>
Contents	<pre> # /etc/bashrc # System wide functions and aliases # Environment stuff goes in /etc/profile # are we an interactive shell? if ["\$PS1"]; then if [-z "\$PROMPT_COMMAND"]; then case \$TERM in xterm*) if [-e /etc/sysconfig/bash-prompt-xterm]; then PROMPT_COMMAND=/etc/sysconfig/bash- prompt-xterm else PROMPT_COMMAND='printf "\033]0;%s@%s:%s\007" "\${USER}" "\${HOSTNAME%%.*}" "\${PWD/#\$HOME/~}"' fi ;; screen) if [-e /etc/sysconfig/bash-prompt-screen]; then PROMPT_COMMAND=/etc/sysconfig/bash- prompt-screen else PROMPT_COMMAND='printf "\033]0;%s@%s:%s\033\\" "\${USER}" "\${HOSTNAME%%.*}" "\${PWD/#\$HOME/~}"' fi ;; *) [-e /etc/sysconfig/bash-prompt-default] && PROMPT_COMMAND=/etc/sysconfig/bash-prompt-default ;; esac fi # Turn on checkwinsize shopt -s checkwinsize ["\$PS1" = "\s-\v\\\\$ "] && PS1="\u@\h \W\\\\$ " fi if ! shopt -q login_shell ; then # We're not a login shell # Need to redefine pathmunge, it get's undefined at the end of /etc/profile pathmunge () { </pre>

	<pre> if ! echo \$PATH /bin/egrep -q "^(:)\$1(\$:)" ; then if ["\$2" = "after"] ; then PATH=\$PATH:\$1 else PATH=\$1:\$PATH fi fi } # By default, we want umask to get set. This sets it for non- login shell. # You could check uidgid reservation validity in # /usr/share/doc/setup-*/uidgid file #if [\$UID -gt 99] && ["`id -gn`" = "`id -un`"]; then # umask 002 #else # umask 022 #fi # Set default permissions to exclude world rwx and group w - ## CRISP Baseline umask 077 # Only display echo's from profile.d scripts if we are no login shell # and interactive - otherwise just process them to set envvars for i in /etc/profile.d/*.sh; do if [-r "\$i"]; then if ["\$PS1"]; then . \$i else . \$i >/dev/null 2>&1 fi fi done unset i unset pathmunge fi # vim:ts=4:sw=4 </pre>
--	--

File	/etc/audit/audit.rules
Settings	Owner: root Group: root Mode: 640
Contents	<pre> # This file contains the auditctl rules that are loaded # whenever the audit daemon is started via the initcripts. # The rules are simply the parameters that would be passed # to auditctl. -D </pre>

	<pre> -b 8192 -f 2 --loginuid-immutable -a always,exit -F arch=b64 -S adjtimex,settimeofday -F key=time- change -a always,exit -F arch=b64 -S clock_settime -F a0=0x0 -F key=time- change -w /etc/localtime -p wa -k time-change -w /etc/group -p wa -k identity -w /etc/passwd -p wa -k identity -w /etc/gshadow -p wa -k identity -w /etc/shadow -p wa -k identity -w /etc/security/opasswd -p wa -k identity -a always,exit -F arch=b64 -S sethostname,setdomainname -F key=system-locale -w /etc/issue -p wa -k system-locale -w /etc/issue.net -p wa -k system-locale -w /etc/hosts -p wa -k system-locale -w /etc/sysconfig/network -p wa -k system-locale -a always,exit -F dir=/etc/NetworkManager/ -F perm=wa -F key=system-locale -a always,exit -F dir=/etc/selinux/ -F perm=wa -F key=MAC-policy -a always,exit -F arch=b64 -S chmod,fchmod,fchmodat -F auid>=1000 -F auid!=4294967295 -F key=perm_mod -a always,exit -F arch=b64 -S chown,fchown,lchown,fchownat -F auid>=1000 -F auid!=4294967295 -F key=perm_mod -a always,exit -F arch=b64 -S setxattr,lsetxattr,fsetxattr,removexattr,lremovexattr,fremovexattr -F auid>=1000 -F auid!=4294967295 -F key=perm_mod -a always,exit -F arch=b32 -S open,creat,truncate,openat,open_by_handle_at -F exit=-EPERM -F auid>=1000 -F auid!=4294967295 -F key=access -a always,exit -F arch=b64 -S open,creat,truncate,creat,openat,open_by_handle_at -F exit=-EACCES -F auid>=1000 -F auid!=4294967295 -F key=access -a always,exit -F arch=b64 -S open,creat,truncate,creat,openat,open_by_handle_at -F exit=-EPERM -F auid>=1000 -F auid!=4294967295 -F key=access -a always,exit -F path=/bin/ping -F perm=x -F auid>=1000 -F auid!=4294967295 -F key=privileged -a always,exit -F arch=b64 -S mount -F auid>=1000 -F auid!=4294967295 -F key=export -a always,exit -F arch=b64 -S unlink,unlinkat,rename,renameat -F auid>=1000 -F auid!=4294967295 -F key=delete -w /etc/sudoers -p wa -k actions -w /etc/sudoers.d/ -p wa -k actions #Make the Audit Configuration Immutable -e 2 </pre>
--	---

File	/etc/cron.allow	
Settings	Owner:	root
	Group:	root
	Mode:	644
Contents	root	

--	--

File	/etc/cron.daily/rpd_update
Settings	Owner: root Group: root Mode: 755
Contents	<pre>#!/bin/bash # rdp_nsupdate v1.6.1 # # This script updates the default names server with the hostname and ip address # including reverse lookup. # It takes an optional parameter of the interface to be used or a config file. # If no parameter is provided and no config file then it will use the first ip address # returned by ifconfig. If you choose to use a config file it must be chmod 600. # (example config file: # # filename: /etc/nsupdate.conf # INTRFC="eth2" # # # end example file) # (example command line: rdp_nsupdate bond0) # 07-19-2013: corrected config file use. # 06/11/2014: added full paths to ensure functionality independent of environment variables # 06/30/2014: moved tmp and log files to /tmp/ and added logging to /var/log/messages # 06/30/2014: added 7 day purge in /tmp to remove older tmp and log files # 12/08/2014: added full path to ifconfig line # 02/25/2015: corrected logging of ntpd sync # TMP_FILE=\$(date +"/tmp/rdp_nsupdate-%Y-%m%d-%H%M%S.tmp") LOG_FILE=\$(date +"/tmp/rdp_nsupdate-%Y-%m%d-%H%M%S.log") CONFIG_FILE=/etc/nsupdate.conf if [\$# -eq 1] then INTRFC=\$1 else if [[-O \$CONFIG_FILE]]; then if [[\$(stat --format %a \$CONFIG_FILE) == 600]]; then . \$CONFIG_FILE fi #original method commented out due to suggestion from Erik #if [[-f \$CONFIG_FILE]]; then # . \$CONFIG_FILE #fi fi #INTRFC=\$1 echo "INTRFC= \$INTRFC " >\$LOG_FILE 2>&1 IPADDR=\$(/sbin/ifconfig \${INTRFC} egrep -o '([0-</pre>

	<pre> 9]{1,3}\.){3}[0-9]{1,3}' sed -n '1p') IFS="." read -ra INARPA <<< "\$IPADDR" echo update add \${HOSTNAME} 86430 a \${IPADDR} >\$TMP_FILE echo send >>\$TMP_FILE echo update add \${INARPA[3]}.\${INARPA[2]}.\${INARPA[1]}.\${INARPA[0]}.in-addr.arpa 86430 ptr \${HOSTNAME} >>\$TMP_FILE echo send >>\$TMP_FILE /usr/bin/nsupdate -v -d \$TMP_FILE >>\$LOG_FILE 2>&1 if [[-x /bin/logger]]; then /bin/logger -i -t rdp_nsupdate -f \$LOG_FILE fi find /tmp/rdp_nsupdate* -type f -mtime +7 -exec rm {} \; 2>/dev/null #rm -f \$TMP_FILE #rm -f \$LOG_FILE /sbin/service ntpd stop >>\$LOG_FILE 2>&1 /usr/sbin/ntpdate -gq >>\$LOG_FILE 2>&1 /sbin/service ntpd start >>\$LOG_FILE 2>&1 </pre>
--	--

File	/etc/csh.login
Settings	Owner: root Group: root Mode: 644
Contents	<pre> # /etc/csh.login # System wide environment and startup programs, for login setup if (\$?PATH) then #do not override user specified PATH else if (\$uid == 0) then setenv PATH "/sbin:/usr/sbin:/usr/local/sbin:/bin:/usr/bin:/usr/local/bin" else setenv PATH "/bin:/usr/bin:/usr/local/bin" endif endif setenv HOSTNAME ` /bin/hostname ` set history=1000 if (! -f \$HOME/.inputrc) then setenv INPUTRC /etc/inputrc endif if (-d /etc/profile.d) then set nonomatch </pre>

	<pre> foreach i (/etc/profile.d/*.csh) if (-r \$i) then if (\$?prompt) then source \$i else source \$i >& /dev/null endif endif end unset i nonomatch endif # Set default permissions to exclude world rwx and group w - ###CRISP Baseline umask 077 </pre>
--	--

File	/etc/inittab
Settings	<pre> Owner: root Group: root Mode: 644 </pre>
Contents	<pre> # inittab is no longer used when using systemd. # # ADDING CONFIGURATION HERE WILL HAVE NO EFFECT ON YOUR SYSTEM. # # Ctrl-Alt-Delete is handled by /usr/lib/systemd/system/ctrl-alt- del.target # # systemd uses 'targets' instead of runlevels. By default, there are two main targets: # # multi-user.target: analogous to runlevel 3 # graphical.target: analogous to runlevel 5 # # To view current default target, run: # systemctl get-default # # To set a default target, run: # systemctl set-default TARGET.target # </pre>

Note: the /etc/inittab file is listed here for completeness, but should not be modified from the file that is deployed by the system at installation.

File	symlink/etc/issue
Settings	
Contents	/etc/issue.net

File	/etc/issue.net
Settings	Owner: root Group: root Mode: 644
Contents	<p>Security Warning!</p> <p>This U.S. government system is intended to be used by authorized VA network users for viewing and retrieving information only except as otherwise explicitly authorized. VA information resides on and transmits through computer systems and networks funded by the VA.</p> <p>All use is considered to be with an understanding and acceptance that there is no reasonable expectation of privacy for any data or transmissions on Government Intranet or Extranet (non-public) networks or systems.,</p> <p>All transactions that occur on this system and all data transmitted through this system are subject to review and action including (but not limited to) monitoring- recording- retrieving- coping- auditing- inspecting- investigating- restricting access- blocking- tracking- disclosing to authorized personnel or any other authorized actions by all authorized VA and law enforcement personnel.,</p> <p>All use of this system constitutes understanding and unconditional acceptance of these terms.,</p> <p>Unauthorized attempts or acts to either (1) access- upload- change- or delete information on this system(2) modify this system (3) deny access to this system or (4) accrue resources for unauthorized use on this system are strictly prohibited., Such attempts or acts are subject to action that may result in criminal civil or administrative penalties</p>

File	/etc/login.defs
Settings	Owner: root Group: root Mode: 644
Contents	<pre># *REQUIRED* # Directory where mailboxes reside, _or_ name of file, relative # to the # home directory. If you _do_ define both, MAIL_DIR takes # precedence. # QMAIL_DIR is for Qmail # #QMAIL_DIR Maildir MAIL_DIR /var/spool/mail #MAIL_FILE .mail # Password aging controls: # # PASS_MAX_DAYS Maximum number of days a password may be # used. # PASS_MIN_DAYS Minimum number of days allowed between</pre>

	<pre> password changes. # PASS_MIN_LEN Minimum acceptable password length. # PASS_WARN_AGE Number of days warning given before a password expires. # # set based on CRISP Baseline PASS_MAX_DAYS 90 PASS_MIN_DAYS 2 PASS_MIN_LEN 8 PASS_WARN_AGE 14 FAIL_DELAY 15 # # Min/max values for automatic uid selection in useradd # UID_MIN 500 UID_MAX 60000 # # Min/max values for automatic gid selection in groupadd # GID_MIN 500 GID_MAX 60000 # # If defined, this command is run when removing a user. # It should remove any at/cron/print jobs etc. owned by # the user to be removed (passed as the first argument). # #USERDEL_CMD /usr/sbin/userdel_local # # If useradd should create home directories for users by default # On RH systems, we do. This option is overridden with the -m flag on # useradd command line. # CREATE_HOME yes # The permission mask is initialized to this value. If not specified, # the permission mask will be initialized to 022. UMASK 077 # This enables userdel to remove user groups if no members exist. # USERGROUPS_ENAB yes # Use MD5 or DES to encrypt password? Red Hat use MD5 by default. # set to SHA512 - CRISP Baseline MD5_CRYPT_ENAB no ENCRYPT_METHOD SHA512 </pre>
--	---

File	/etc/logrotate.conf
Settings	Owner: root Group: root Mode: 644
Contents	<pre># see "man logrotate" for details # rotate log files weekly weekly # keep 4 weeks' worth of backlogs rotate 4 # create new (empty) log files after rotating old ones create # uncomment this if you want your log files compressed compress # RPM packages drop log rotation information into this directory include /etc/logrotate.d # no packages own wtmp -- we'll rotate them here /var/log/wtmp { monthly minsize 1M create 0664 root utmp rotate 1 } /var/log/btmp { missingok monthly minsize 1M create 0600 root utmp rotate 1 } # system-specific logs may be also be configured here.</pre>

File	/etc/logrotate.d/syslog
Settings	Owner: root Group: root Mode: 644
Contents	<pre>/var/log/messages /var/log/secure /var/log/maillog /var/log/spooler /var/log/boot.log /var/log/cron /var/log/sudo.log { sharedscripts postrotate /bin/kill -HUP `cat /var/run/syslogd.pid 2> /dev/null` 2></pre>

	<pre> /dev/null true /bin/kill -HUP `cat /var/run/rsyslogd.pid 2> /dev/null` 2> /dev/null true endscrip } </pre>
--	--

File	/etc/modprobe.d/crisp.conf
Settings	<pre> Owner: root Group: root Mode: 644 </pre>
Contents	<pre> #Disable jffs2 install jffs2 /bin/true #Disable hfs install hfs /bin/true #Disable HFS Plus install hfsplus /bin/true #Disable squashfs install squashfs /bin/true #Disable udf install udf /bin/true #Disable cramfs install cramfs /bin/true #Disable freexfw install freevxfs /bin/true #disable bluetooth install net-pf-31 /bin/false #disable sctp install sctp /bin/false #disable dccp install dccp /bin/false #Disable usb storage install usb-storage /bin/false </pre>

File	/etc/ntp.conf
Settings	<pre> Owner: root Group: root Mode: 644 </pre>
Contents	<pre> # For more information about this file, see the man pages # ntp.conf(5), ntp_acc(5), ntp_auth(5), ntp_clock(5), ntp_misc(5), ntp_mon(5). driftfile /var/lib/ntp/drift # Permit time synchronization with our time source, but do not # permit the source to query or modify the service on this system. restrict default kod nomodify notrap nopeer noquery restrict -6 default kod nomodify notrap nopeer noquery </pre>

	<pre> # Permit all access over the loopback interface. This could # be tightened as well, but to do so would affect some of # the administrative functions. restrict 127.0.0.1 restrict -6 ::1 server ntp.va.gov iburst server ntp.va.gov iburst server ntp.va.gov iburst # Key file containing the keys and key identifiers used when operating # with symmetric key cryptography. keys /etc/ntp/keys </pre>
--	--

File	/etc/ntp/step-tickers
Settings	<pre> Owner: root Group: root Mode: 644 </pre>
Contents	<pre> ntp.va.gov ntp3.va.gov ntp4.va.gov ntp5.va.gov ntp1.va.gov ntp6.va.gov ntp2.va.gov 127.127.1.0 </pre>

File	/etc/pam.d/system-auth-va
Settings	<pre> Owner: root Group: root Mode: 644 </pre>
Contents	<pre> #%PAM-1.0 # RHEL 6 Baseline configuration auth required pam_env.so auth [default=die] pam_faillock.so authfail deny=3 unlock_time=900 fail_interval=900 auth required pam_faillock.so authsucc deny=3 unlock_time=900 fail_interval=900 auth sufficient pam_unix.so nullok try_first_pass auth requisite pam_succeed_if.so uid >= 500 quiet auth required pam_deny.so account required pam_unix.so </pre>

	account	sufficient	pam_localuser.so
	account	sufficient	pam_succeed_if.so uid < 500 quiet
	account	required	pam_permit.so
	password	requisite	pam_cracklib.so try_first_pass retry=5
	minclass=3	minlen=8	difok=4
	password	sufficient	pam_unix.so sha512 shadow nullok
	try_first_pass	use_authtok	remember=5
	password	required	pam_deny.so
	session	optional	pam_keyinit.so revoke
	session	required	pam_limits.so
	session	[success=1	default=ignore] pam_succeed_if.so service
	in crond	quiet	use_uid
	session	required	pam_unix.so

File	/etc/pam.d/su		
Settings	Owner:	root	
	Group:	root	
	Mode:	644	
Contents	<pre> #%PAM-1.0 auth sufficient pam_rootok.so # Uncomment the following line to implicitly trust users in the # "wheel" group. #auth sufficient pam_wheel.so trust use_uid # Require wheel group for su - CRISP Baseline # RHEL 7 Baseline auth required pam_wheel.so use_uid auth include system-auth account sufficient pam_succeed_if.so uid = 0 use_uid quiet account include system-auth password include system-auth session include system-auth session optional pam_xauth.so </pre>		

File	/etc/profile.d/halt.sh		
Settings	Owner:	root	
	Group:	root	
	Mode:	644	
Contents	alias halt='echo "Do not do that here"' 2>/dev/null		

File	/etc/profile.d/tmout.csh		
Settings	Owner:	root	
	Group:	root	
	Mode:	644	
Contents	set -r autologout 15		

File	/etc/profile.d/tmout.sh
Settings	Owner: root Group: root Mode: 644
Contents	TMOUT=900 readonly TMOUT export TMOUT

File	/etc/securetty
Settings	Owner: root Group: root Mode: 600
Contents	console tty1 tty2 tty3 tty4 tty5 tty6

File	/etc/sysconfig/prelink
Settings	Owner: root Group: root Mode: 600
Contents	PRELINKING=no

File	/etc/sysctl.d/50-crisp.conf
Settings	Owner: root Group: root Mode: 644
Contents	# Kernel sysctl configuration file for Red Hat Linux # # For binary values, 0 is disabled, 1 is enabled. # Controls the System Request debugging functionality of # the kernel kernel.sysrq = 0 # Controls whether core dumps will append the PID to the core # filename. # Useful for debugging multi-threaded applications. kernel.core_uses_pid = 1 # CRISP Baseline

	<pre> # ExecShield describes kernel features that provide protection #against exploitation of memory corruption errors such as #buffer overflows kernel.exec-shield=1 # Enable Randomized Layout of Virtual Address Space kernel.randomize_va_space = 2 #disable core dumps for SUID programs fs.suid_dumpable = 0 # Controls IP packet forwarding net.ipv4.ip_forward = 0 # Controls source route verification net.ipv4.conf.default.rp_filter = 1 net.ipv4.conf.all.rp_filter = 1 # Do not accept source routing net.ipv4.conf.default.accept_source_route = 0 net.ipv4.conf.all.accept_source_route = 0 # Protect against SYN floods net.ipv4.tcp_max_syn_backlog = 4096 net.ipv4.tcp_syncookies = 1 # Do not participate in SMURF attacks net.ipv4.icmp_echo_ignore_broadcasts = 1 # Do not be overly verbose in logging net.ipv4.icmp_ignore_bogus_error_responses = 1 # Do not accept or send redirects net.ipv4.conf.default.accept_redirects =0 net.ipv4.conf.all.accept_redirects = 0 net.ipv4.conf.default.secure_redirects = 0 net.ipv4.conf.all.secure_redirects = 0 net.ipv4.conf.default.send_redirects = 0 net.ipv4.conf.all.send_redirects = 0 </pre>
--	---

File	/etc/security/limits.conf			
Settings	Owner:	root		
	Group:	root		
	Mode:	644		
Contents	# /etc/security/limits.conf			
	#			
	#*	soft	core	0
	#*	hard	rss	10000
	#@student	hard	nproc	20
	#@faculty	soft	nproc	20
	#@faculty	hard	nproc	50
	#ftp	hard	nproc	0

	#@student	-	maxlogins	4
--	-----------	---	-----------	---

File	/etc/security/limits.d/50-crisp.conf			
Settings	Owner:	root		
	Group:	root		
	Mode:	644		
Contents	# CRISP Baseline			
	*	hard	core	0
	*	hard	fsize	102400
	*	hard	nproc	150
	root	-	fsize	-1
	root	-	nproc	-1
	# End of file			

File	/etc/selinux/config
Settings	Owner: root Group: root Mode: 644
Contents	<pre># This file controls the state of SELinux on the system. # SELINUX= can take one of these three values: # enforcing - SELinux security policy is enforced. # permissive - SELinux prints warnings instead of # enforcing. # disabled - SELinux is fully disabled. SELINUX=permissive # SELINUXTYPE= type of policy in use. Possible values are: # targeted - Only targeted network daemons are protected. # strict - Full SELinux protection. SELINUXTYPE=targeted # SETLOCALDEFS= Check local definition changes SETLOCALDEFS=0</pre>

File	/etc/ssh/sshd_config		
Settings	Owner:	root	
	Group:	root	
	Mode:	644	
Contents	<pre># # \$OpenBSD: sshd_config,v 1.80 2008/07/02 02:24:18 djm Exp \$ # This is the sshd server system-wide configuration file. See # sshd_config(5) for more information. # This sshd was compiled with PATH=/usr/local/bin:/bin:/usr/bin</pre>		

	<pre> # The strategy used for options in the default sshd_config shipped with # OpenSSH is to specify options with their default value where # possible, but leave them commented. Uncommented options change a # default value. #Port 22 #AddressFamily any #ListenAddress 0.0.0.0 #ListenAddress :: # Disable legacy (protocol version 1) support in the server for new # installations. In future the default will change to require explicit # activation of protocol 1 Protocol 2 #Use only fips certified ciphers Ciphers aes128-ctr,aes192-ctr,aes256-ctr,aes128-cbc,3des- cbc,aes192-cbc,aes256-cbc #Disable MD5 MACs (message authentication code) algorithms MACs hmac-sha1,hmac-ripemd160 # HostKey for protocol version 1 #HostKey /etc/ssh/ssh_host_key # HostKeys for protocol version 2 #HostKey /etc/ssh/ssh_host_rsa_key #HostKey /etc/ssh/ssh_host_dsa_key # Lifetime and size of ephemeral version 1 server key #KeyRegenerationInterval 1h #ServerKeyBits 1024 # Logging # obsoletes QuietMode and FascistLogging #SyslogFacility AUTH SyslogFacility AUTHPRIV #LogLevel INFO # Authentication: #LoginGraceTime 2m PermitRootLogin no #StrictModes yes #MaxAuthTries 6 #MaxSessions 10 #RSAAuthentication yes #PubkeyAuthentication yes #AuthorizedKeysFile .ssh/authorized_keys #AuthorizedKeysCommand none #AuthorizedKeysCommandRunAs nobody </pre>
--	---

	<pre> # For this to work you will also need host keys in /etc/ssh/ssh_known_hosts #RhostsRSAAuthentication no # similar for protocol version 2 #HostbasedAuthentication no # Change to yes if you don't trust ~/.ssh/known_hosts for # RhostsRSAAuthentication and HostbasedAuthentication #IgnoreUserKnownHosts no # Don't read the user's ~/.rhosts and ~/.shosts files #IgnoreRhosts yes # To disable tunneled clear text passwords, change to no here! #PasswordAuthentication yes #PermitEmptyPasswords no PasswordAuthentication yes # Change to no to disable s/key passwords #ChallengeResponseAuthentication yes ChallengeResponseAuthentication no # Kerberos options #KerberosAuthentication no #KerberosOrLocalPasswd yes #KerberosTicketCleanup yes #KerberosGetAFSToken no #KerberosUseKuserok yes # GSSAPI options #GSSAPIAuthentication no GSSAPIAuthentication yes #GSSAPICleanupCredentials yes GSSAPICleanupCredentials yes #GSSAPIStrictAcceptorCheck yes #GSSAPIKeyExchange no # Set this to 'yes' to enable PAM authentication, account processing, # and session processing. If this is enabled, PAM authentication will # be allowed through the ChallengeResponseAuthentication and # PasswordAuthentication. Depending on your PAM configuration, # PAM authentication via ChallengeResponseAuthentication may bypass # the setting of "PermitRootLogin without-password". # If you just want the PAM account and session checks to run without # PAM authentication, then enable this but set PasswordAuthentication # and ChallengeResponseAuthentication to 'no'. #UsePAM no UsePAM yes # Accept locale-related environment variables AcceptEnv LANG LC_CTYPE LC_NUMERIC LC_TIME LC_COLLATE LC_MONETARY LC_MESSAGES AcceptEnv LC_PAPER LC_NAME LC_ADDRESS LC_TELEPHONE LC_MEASUREMENT AcceptEnv LC_IDENTIFICATION LC_ALL LANGUAGE </pre>
--	---

	<pre> AcceptEnv XMODIFIERS #AllowAgentForwarding yes #AllowTcpForwarding yes #GatewayPorts no #X11Forwarding no X11Forwarding yes #X11DisplayOffset 10 #X11UseLocalhost yes #PrintMotd yes #PrintLastLog yes #TCPKeepAlive yes #UseLogin no #UsePrivilegeSeparation yes #PermitUserEnvironment no #Compression delayed #ClientAliveInterval 0 #ClientAliveCountMax 3 #ShowPatchLevel no #UseDNS yes UseDNS no #PidFile /var/run/sshd.pid #MaxStartups 10 #PermitTunnel no #ChrootDirectory none #default banner path Banner /etc/issue.net # override default of no subsystems Subsystem sftp /usr/libexec/openssh/sftp-server # Example of overriding settings on a per-user basis #Match User anoncvs # X11Forwarding no # AllowTcpForwarding no # ForceCommand cvs server # </pre>
--	--

File	/etc/sudoers
Settings	<pre> Owner: root Group: root Mode: 440 </pre>
Contents	<pre> ## Networking Cmnd_Alias NETWORKING = /sbin/route, /sbin/ifconfig, /bin/ping, /sbin/dhclient, /usr/bin/net, /sbin/iptables, /usr/bin/rfcomm, /usr/bin/wvdial, /sbin/iwconfig, /sbin/mii-tool ## Installation and management of software Cmnd_Alias SOFTWARE = /bin/rpm, /usr/bin/up2date, /usr/bin/yum ## Services Cmnd_Alias SERVICES = /sbin/service, /sbin/chkconfig </pre>

```

## Updating the locate database
Cmnd_Alias LOCATE = /usr/sbin/updatedb

## Storage
Cmnd_Alias STORAGE = /sbin/fdisk, /sbin/sfdisk, /sbin/parted,
/sbin/partprobe, /bin/mount, /bin/umount

## Delegating permissions
Cmnd_Alias DELEGATING = /usr/sbin/visudo, /bin/chown, /bin/chmod,
/bin/chgrp

## Processes
Cmnd_Alias PROCESSES = /bin/nice, /bin/kill, /usr/bin/kill,
/usr/bin/killall

## Drivers
Cmnd_Alias DRIVERS = /sbin/modprobe

# Defaults specification
# Disable "ssh hostname sudo <cmd>", because it will show the
password in clear.
#         You have to run "ssh -t hostname sudo <cmd>".
#
Defaults        requiretty

Defaults        env_reset
Defaults        env_keep = "COLORS DISPLAY HOSTNAME HISTSIZE INPUTRC
KDEDIR \

                                LS_COLORS MAIL PS1 PS2 QTDIR USERNAME \
                                LANG LC_ADDRESS LC_CTYPE LC_COLLATE

LC_IDENTIFICATION \
                                LC_MEASUREMENT LC_MESSAGES LC_MONETARY
LC_NAME LC_NUMERIC \
                                LC_PAPER LC_TELEPHONE LC_TIME LC_ALL
LANGUAGE LINGUAS \
                                _XKB_CHARSET XAUTHORITY"

# Log to /var/log/sudo.log - CRISP Baseline
Defaults        logfile=/var/log/sudo.log,loglinelen=0

##
## Allow root to run any commands anywhere
root    ALL=(ALL)        ALL

## Allows people in group wheel to run all commands
%wheel   ALL=(ALL)        ALL

## Same thing without a password
# %wheel  ALL=(ALL)        NOPASSWD: ALL

## Read drop-in files from /etc/sudoers.d (the # here does not
mean a comment)
#includedir /etc/sudoers.d

```

File	/etc/sysctl.d/ipv6.conf
Settings	Owner: root Group: root Mode: 644
Contents	net.ipv6.conf.all.disable_ipv6 = 1 net.ipv6.conf.default.disable_ipv6 = 1 #Disable accepting IPv6 router advertisements net.ipv6.conf.default.accept_ra = 0 # Disable Accepting IPv6 Redirects net.ipv6.conf.default.accept_redirects = 0